

Simultaneously Structured Models with Application to Sparse and Low-rank Matrices

Samet Oymak^c, Amin Jalali^w,
Maryam Fazel^w, Yonina C. Eldar^t, Babak Hassibi^c

^c California Institute of Technology

^w University of Washington

^t Technion

February 12, 2013

Abstract

The topic of recovery of a structured model given a small number of linear observations has been well-studied in recent years. Examples include recovering sparse or group-sparse vectors, low-rank matrices, and the sum of sparse and low-rank matrices, among others. In various applications in signal processing and machine learning, the model of interest is known to be structured in *several* ways at the same time, for example, a matrix that is simultaneously sparse and low-rank. An important application is the sparse phase retrieval problem, where the goal is to recover a sparse signal from phaseless measurements. In machine learning, the problem comes up when combining several regularizers that each promote a certain desired structure.

Often penalties (norms) that promote each individual structure are known and yield an order-wise optimal number of measurements (e.g., ℓ_1 norm for sparsity, nuclear norm for matrix rank), so it is reasonable to minimize a combination of such norms. We show that, surprisingly, if we use multi-objective optimization with the individual norms, then we can do no better, order-wise, than an algorithm that exploits only one of the several structures. This result suggests that to fully exploit the multiple structures, we need an entirely new convex relaxation, i.e., not one that is a function of the convex relaxations used for each structure. We then specialize our results to the case of sparse and low-rank matrices. We show that a nonconvex formulation of the problem can recover the model from very few measurements, on the order of the degrees of freedom of the matrix, whereas the convex problem obtained from a combination of the ℓ_1 and nuclear norms requires many more measurements. This proves an order-wise gap between the performance of the convex and nonconvex recovery problems in this case.

Keywords. compressed sensing, convex relaxation, regularization, performance bounds

1 Introduction

Recovery of a structured model (signal) given a small number of linear observations has been the focus of many studies recently. Examples include recovering sparse or group-sparse vectors (which gave rise to the area of compressed sensing [1, 2, 3], low-rank matrices [4, 5], and the sum of sparse and low-rank matrices [6, 7], among others. More generally, the recovery of a signal that can be expressed as the sum of a few atoms out of a appropriate atomic set, has been studied in [8]. Canonical questions that have guided research in this area include: How many generic linear measurements are enough to recover the model by any means? How many measurements are enough for a tractable approach, e.g., solving a convex optimization problem? In the statistics literature, these questions are posed in terms of “sample complexity” and error rates for estimators minimizing the sum of a quadratic loss function and a regularizer that reflects the desired structure [9].

In practice, there are many cases where the model of interest is known to be structured in *several* ways at the same time. We then seek a signal that lies in the intersection of several sets defining the individual structures (in a sense that we will make precise later). Such recovery problems arise often in applications, for example in signal processing (Section 1.2) as well as statistical learning. The most common convex regularizer (norm penalty), used to promote all structures together, is a linear combination of well-known regularizers for each structure. However, there is currently no general analysis and understanding of how well such regularization performs, in terms of the number of observations required for successful recovery of the desired model. This paper addresses this ubiquitous yet unexplored problem; i.e., the recovery of *simultaneously structured models*.

An example of a simultaneously structured model is a matrix that is *simultaneously sparse and low-rank*. In this case, one would like to come up with algorithms that exploit both types of structures to minimize the number of measurements required for recovery. An $n \times n$ matrix with rank $r \ll n$ can be described by $O(rn)$ parameters, and can be recovered using $O(rn)$ generic measurements via nuclear norm minimization [4, 10]. On the other hand, a block-sparse matrix with a $k \times k$ nonzero block where $k \ll n$ can be described by k^2 parameters and can be recovered with $O(k^2 \log(n/k))$ generic measurements using ℓ_1 minimization. However, a matrix that is *both* rank r and block-sparse can be described by $O(rk)$ parameters. The question is whether we can exploit this joint structure to efficiently recover such a matrix with $O(rk)$ measurements.

In this paper we answer this question in the negative, in the following sense: if we use multi-objective optimization with the ℓ_1 norm and the nuclear norm (used for sparse signals and low rank matrices, respectively), then the number of measurements required is lower bounded by $O(\min\{k^2, rn\})$. In other words, we need at least this number of observations for the desired signal to lie on the Pareto optimal front traced by the objectives, the ℓ_1 norm and the nuclear norm. This means we can do *no better than an algorithm that exploits only one of the two structures*.

We introduce a framework to express general simultaneous structures, and as our main result, we prove that the same phenomenon happens for a general set of structures under reasonable assumptions on the norm penalties used. These assumptions hold in many typical cases of interest, such as combinations of sparse, group-sparse, and low-rank structures. The measurements we consider are *generic* measurements; we focus on random Gaussian measurement matrices, with independent and identically distributed entries. This gives an open, dense subset of the set of all m -measurement matrices, hence justifying the term “generic”.

Table 1 summarizes known results on recovery of some common structured models, along with a result of this paper specialized to the problem of low-rank and sparse matrix recovery. The first column gives the number of parameters needed to describe the model (often referred to as its ‘degrees of freedom’), the second and third columns show how many generic measurements are needed for successful recovery. In using ‘nonconvex recovery’, we assume we are able to find the global minimum of a nonconvex problem. This is clearly intractable in general, and not a practical recovery method—we consider it as a benchmark for theoretical comparison with the (tractable) convex relaxation, in order to determine how powerful the relaxation is.

The first and second rows are the results on k sparse vectors in \mathbb{R}^n and rank r matrices in $\mathbb{R}^{n \times n}$ respectively, [11, 10]. The third row considers the recovery of “low-rank plus sparse” matrices. Consider a matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$ that can be decomposed as $\mathbf{X} = \mathbf{X}_L + \mathbf{X}_S$ where \mathbf{X}_L is a rank r matrix and \mathbf{X}_S is a matrix with only k nonzero entries. The degrees of freedom of \mathbf{X} is $O(rn + k)$. Minimizing the infimal convolution of ℓ_1 norm and nuclear norm, i.e., $f(\mathbf{X}) = \min_{\mathbf{Y}} \|\mathbf{Y}\|_* + \lambda \|\mathbf{X} - \mathbf{Y}\|_1$ subject to random Gaussian measurements on \mathbf{X} , gives a convex approach for recovering \mathbf{X} . It has been shown that under reasonable incoherence assumptions, \mathbf{X} can be recovered uniquely from $O((rn + k) \log^2 n)$ measurements which is suboptimal only by a logarithmic factor [12]. Finally, the last row in Table 1 shows one of the results in this paper. Let $\mathbf{X} \in \mathbb{R}^{n \times n}$ be a rank r matrix whose entries are zero outside a $k_1 \times k_2$ submatrix. The degrees of freedom of \mathbf{X} is $O((k_1 + k_2)r)$. We consider both convex and non-convex programs for the recovery of this type of matrices. The nonconvex method involves minimizing the number of nonzero rows, columns and rank of the matrix jointly, as discussed in Section 3.2. As it will be shown later, $O((k_1 + k_2)r \log n)$ measurements suffices for this program to successfully recover the original matrix. The convex method minimizes any convex combination of the individual structure-inducing norms, namely the nuclear norm and the $\ell_{1,2}$ norms of the rows

Model	Degrees of Freedom	Nonconvex recovery	Convex recovery
Sparse vectors	k	$O(k)$	$O(k \log \frac{n}{k})$
Low rank matrices	$r(2n - r)$	$O(rn)$	$O(rn)$
Low rank plus sparse	$O(rn + k)$	not analyzed	$O((rn + k) \log^2 n)$
Low rank and sparse	$O(r(k_1 + k_2))$	$O(r(k_1 + k_2) \log n)$	$\Omega(rn)$

Table 1: Summary of results in recovery of structured signals. This paper shows a gap between the performance of convex and nonconvex recovery programs for simultaneously structured matrices (last row).

and columns, [13, 14]. We show that with high probability this program cannot recover the original matrix with fewer than $\Omega(rn)$ measurements. In summary, while nonconvex method performs slightly suboptimal, the convex method performs poorly as the amount of measurements scales with n rather than $k_1 + k_2$.

1.1 Contributions

This paper describes a general framework for analyzing the recovery of models that have more than one structure, by combining penalty functions corresponding to each structure. The framework proposed includes special cases that are of interest in their own right, e.g., sparse and low-rank matrix recovery. Our contributions can be summarized as follows.

Poor performance of convex relaxations. We consider a model with several structures with the assumption that all structure-inducing norms are decomposable at the true input signal \mathbf{x}_0 (see Section 2). For recovery, we consider a multi-objective optimization problem to minimize the individual norms simultaneously. Using Pareto optimality, we know that minimizing a weighted sum of the norms and varying the weights traces out all points of the Pareto-optimal front (i.e., the trade-off surface, Section 2). We obtain a lower bound on the number of measurements, that holds no matter what the weights are and no matter what function is used to trace the points on the Pareto-optimal front. A sketch of our main result is as follows.

Given a model (signal) \mathbf{x}_0 with τ simultaneous structures satisfying certain conditions, the number of generic measurements required for recovery with high probability using any linear combination of the individual norms satisfies the lower bound

$$m \geq c d_{\min} = c \min_{i=1, \dots, \tau} d_i$$

where d_i is approximately on the order of the number of measurements required if minimizing the i th norm only. The constant c depends on the individual norms, as well as the relative geometry of their norm balls at \mathbf{x}_0 .

With d_{\min} as the bottleneck, this result implies that the combination of norms can perform no better than using only one of the norms, even though the target model is tightly constrained and has a very small degree of freedom.

Incorporating general cone constraints. Our results incorporate side information on \mathbf{x}_0 , expressed as convex cone constraints. This additional information about the signal helps in recovery; however, quantifying how much the cone constraints can help is not trivial. Our analysis explicitly determines the role of the cone constraint: Geometric properties of the cone such as its Gaussian width determine the constant factors in the bound on the number of measurements.

Sparse and Low-rank matrix recovery: illustrating a gap. As a special case, we consider the recovery of simultaneously sparse and low-rank matrices and prove that there is a significant gap between the performance of convex and non-convex recovery programs. This gap is surprising when one considers similar results in low-dimensional model recovery discussed above in Table 1.

1.2 Some applications

Here we survey several applications of the sparse and low-rank matrix recovery problem, as well as existing results specific to these applications. A related problem is the Sparse Principal Component Analysis problem, mentioned in Section 7. Other applications include Collaborative Hierarchical Sparse Modeling [15], where sparsity is considered within the non-zero blocks in a block-sparse vector and the recovery of hyperspectral images where we aim to recover a simultaneously block sparse and low rank matrix from compressed observations, [16].

Sparse signal recovery from quadratic measurements. Sparsity has long been exploited in signal processing, applied mathematics, statistics and computer science for tasks such as compression, denoising, model selection, image processing and more. Despite the great interest in exploiting sparsity in various applications, most of the work to date has focused on recovering sparse or low rank data from linear measurements. Recently, the basic sparse recovery problem has been generalized to the case in which the measurements are given by general nonlinear transforms of the unknown input, [18]. A special case of this more general setting is quadratic compressed sensing introduced in [17] in which the goal is to recover a sparse vector \mathbf{x} from quadratic measurements $b_i = \mathbf{x}\mathbf{A}_i\mathbf{x}^T$. This problem can be linearized by *lifting*, where we wish to recover a “low rank and sparse” matrix $\mathbf{X} = \mathbf{x}\mathbf{x}^T$ subject to measurements $b_i = \langle \mathbf{A}_i, \mathbf{X} \rangle$.

Sparse recovery problems from quadratic measurements arise in a variety of different problems in optics. One example is sub-wavelength optical imaging [19, 17] in which the goal is to recover a sparse image from its far-field measurements, where due to the laws of physics the relationship between the (clean) measurement and the unknown image is quadratic. In [17] the quadratic relationship is a result of using partially-incoherent light. The quadratic behavior of the measurements in [19] arises from coherent diffractive imaging in which the image is recovered from its intensity pattern. Under an appropriate experimental setup, this problem amounts to reconstruction of a sparse signal from the magnitude of its Fourier transform.

Sparse phase retrieval. Quadratic measurements appear in phase retrieval problems, in which a signal is to be recovered from the magnitude of its measurements $b_i = |\mathbf{a}_i^T \mathbf{x}|$, where each measurement is a linear transform of the input $\mathbf{x} \in \mathbb{R}^n$ and \mathbf{a}_i are arbitrary possibly complex-valued measurement vectors. Phase retrieval is of great interest in many applications such as optical imaging [20, 21], crystallography [22], and more [23].

The problem becomes linear when \mathbf{x} is *lifted* and we consider the recovery of $\mathbf{X} = \mathbf{x}\mathbf{x}^T$ where each measurement takes the form $b_i^2 = \langle \mathbf{a}_i\mathbf{a}_i^T, \mathbf{X} \rangle$. In [17], an algorithm was developed to treat phase retrieval problems with sparse \mathbf{x} based on a semidefinite relaxation, and low-rank matrix recovery combined with a row-sparsity constraint on the resulting matrix. More recent works also proposed the use of semidefinite relaxation together with sparsity constraints for phase retrieval [24, 25, 26, 27]. An alternative algorithm was recently designed in [28] based on a greedy search. In [26], the authors also consider sparse signal recovery based on combinatorial and probabilistic approaches and give uniqueness results under certain conditions. Stable uniqueness in phase retrieval problems is studied in [29]. The results of [30, 31, 32, 33] applies to general (non-sparse) signals where in some cases *masked* versions of the signal are required.

The problem of recovering a signal from the magnitude of its Fourier transform has been studied extensively in the literature. Many methods have been developed for phase recovery [23] which often rely on prior information about the signal, such as positivity or support constraints. One of the most popular techniques is based on alternating projections, where the current signal estimate is transformed back and forth between the object and the Fourier domains. The prior information and observations are used in each domain in order to form the next estimate. Two of the main approaches of this type are Gerchberg-Saxton [35] and Fienup [34]. In general, these methods are not guaranteed to converge, and often require careful parameter selection and sufficient signal constraints in order to provide a reasonable result. Approaches based on semidefinite relaxation or the recently proposed greedy methods appear to lead to far superior recovery and convergence results.

1.3 Outline of the paper

The paper is structured as follows. Background, definitions, and assumptions are given in Section 2. An overview of the main results is provided in Section 3. The proofs of the general results are presented in Section 4. The proofs for the special case of simultaneously sparse and low-rank matrices are discussed in Section 5, where we compare corollaries of the general results with the results on non-convex recovery approaches and illustrate a gap. Numerical simulations in Section 6 empirically support the results on sparse and low-rank matrices. Future directions of research and discussion of results are in Section 7.

2 Problem Setup

We begin by reviewing some basic definitions. For a vector $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|$ denotes a general norm and $\|\mathbf{x}\|^* = \sup_{\|\mathbf{z}\| \leq 1} \langle \mathbf{x}, \mathbf{z} \rangle$ is the corresponding dual norm. A subgradient of the norm $\|\cdot\|$ at \mathbf{x} is a vector \mathbf{g} for which $\|\mathbf{z}\| \geq \|\mathbf{x}\| + \langle \mathbf{g}, \mathbf{z} - \mathbf{x} \rangle$ holds for any \mathbf{z} . The set of all subgradients is called the subdifferential, denoted by $\partial\|\mathbf{x}\|$. We consider finite dimensional spaces and convex functions, so the subdifferential $\partial\|\mathbf{x}\|$ is always a compact convex set. For a convex set M and point \mathbf{x} , we define the projection operator as

$$\mathcal{P}_M(\mathbf{x}) = \arg \min_{\mathbf{u} \in M} \|\mathbf{x} - \mathbf{u}\|_2.$$

For a subspace M , denote its orthogonal complement by M^\perp . The set of $n \times n$ positive semidefinite (PSD) and symmetric matrices will be denoted by \mathbb{S}_+^n and \mathbb{S}^n respectively. Given a cone \mathcal{C} , denote its dual by \mathcal{C}^* , defined as

$$\mathcal{C}^* = \{\mathbf{z} \mid \langle \mathbf{z}, \mathbf{v} \rangle \geq 0 \text{ for all } \mathbf{v} \in \mathcal{C}\}. \quad (2.1)$$

The polar cone is denoted by \mathcal{C}° and is given by $\mathcal{C}^\circ = -\mathcal{C}^*$.

When considering simultaneously structured signals, we restrict our attention to structures associated with norms that satisfy a *decomposability* property, defined next. The definition is sufficiently general to cover popular structure-inducing norms. Examples include the ℓ_1 norm which induces vector sparsity, the $\ell_{1,2}$ norm which induces column sparsity in a matrix, and the nuclear norm which promotes a low-rank matrix. The nuclear norm gives the summation of the singular values of a matrix and it will be denoted as $\|\cdot\|_\star$.

Definition 2.1 (Decomposable Norm) *A norm $\|\cdot\|$ is decomposable at $\mathbf{x} \in \mathbb{R}^n$ if there exist a subspace $T \subset \mathbb{R}^n$ and a vector $\mathbf{e} \in T$ such that the subdifferential at \mathbf{x} has the form*

$$\partial\|\mathbf{x}\| = \{\mathbf{z} \in \mathbb{R}^n : \mathcal{P}_T(\mathbf{z}) = \mathbf{e}, \|\mathcal{P}_{T^\perp}(\mathbf{z})\|^* \leq 1\}, \quad (2.2)$$

and for all $\mathbf{s} \in T^\perp$ we have

$$\|\mathbf{s}\| = \sup_{\mathbf{z} \in T^\perp, \|\mathbf{z}\|^* \leq 1} \langle \mathbf{s}, \mathbf{z} \rangle. \quad (2.3)$$

We refer to T as the support and \mathbf{e} as the sign vector of \mathbf{x} with respect to $\|\cdot\|$.

Definition 2.1 is used in [36], and is closely related to the one given in [12] where the authors assume *non-expansiveness* of \mathcal{P}_{T^\perp} with respect to dual norm $\|\cdot\|^*$ instead of (2.3); that is, $\|\mathcal{P}_{T^\perp}(\mathbf{x})\|^* \leq \|\mathbf{x}\|^*$ holds for all $\mathbf{x} \in \mathbb{R}^n$ (this condition implies (2.3)).

The following lemma gives a useful relation between \mathbf{x} and the corresponding sign and support.

Lemma 2.1 *Let $\|\cdot\|, \mathbf{x}, \mathbf{e}$ and T be as in Definition 2.1. We have that $\mathbf{x} \in T$ and $\langle \mathbf{x}, \mathbf{e} \rangle > 0$.*

Proof. Using the definition of subgradient of a norm, it can be shown that for any $\mathbf{g} \in \partial\|\mathbf{x}\|$ we have $\langle \mathbf{x}, \mathbf{g} \rangle = \|\mathbf{x}\|$, [37]. From (2.2), we know that $\mathbf{e} \in \partial\|\mathbf{x}\|$, hence $\langle \mathbf{x}, \mathbf{e} \rangle = \|\mathbf{x}\| > 0$. Now, using (2.3), choose $\mathbf{z} \in T^\perp$ such that $\langle \mathbf{x}, \mathbf{z} \rangle = \|\mathcal{P}_{T^\perp}(\mathbf{x})\|$ and $\|\mathbf{z}\|^* \leq 1$. Then $\mathbf{e} + \mathbf{z} \in \partial\|\mathbf{x}\|$ and

$$\|\mathbf{x}\| = \langle \mathbf{x}, \mathbf{e} + \mathbf{z} \rangle = \|\mathbf{x}\| + \|\mathcal{P}_{T^\perp}(\mathbf{x})\|, \quad (2.4)$$

implying $\|\mathcal{P}_{T^\perp}(\mathbf{x})\| = 0$ or $\mathcal{P}_{T^\perp}(\mathbf{x}) = 0$, which means $\mathbf{x} \in T$. ■

To give some intuition for Definition 2.1, we review examples of norms that arise when considering simultaneously sparse and low rank matrices. For a matrix $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$, let $\mathbf{X}_{i,j}$, $\mathbf{X}_{i,\cdot}$ and $\mathbf{X}_{\cdot,j}$ denote its (i, j) entry, i th row and j th column respectively.

Lemma 2.2 (see [36]) $\ell_1, \ell_{1,2}$ and $\|\cdot\|_\star$ are decomposable as follows.

- **ℓ_1 norm** is decomposable at every $\mathbf{x} \in \mathbb{R}^n$, with sign $\mathbf{e} = \text{sgn}(\mathbf{x})$, and support as

$$T = \text{supp}(\mathbf{x}) = \{\mathbf{y} \in \mathbb{R}^n : \mathbf{x}_i = 0 \Rightarrow \mathbf{y}_i = 0 \text{ for } i = 1, \dots, n\}.$$

- **$\ell_{1,2}$ norm** is decomposable at every $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$. The support is

$$T = \{\mathbf{Y} \in \mathbb{R}^{n_1 \times n_2} : \mathbf{X}_{\cdot,i} = \mathbf{0} \Rightarrow \mathbf{Y}_{\cdot,i} = \mathbf{0} \text{ for } i = 1, \dots, n_2\},$$

and the sign vector $\mathbf{e} \in \mathbb{R}^{n_1 \times n_2}$ is obtained by normalizing the columns of \mathbf{X} present in the support, $\mathbf{e}_{\cdot,j} = \frac{\mathbf{X}_{\cdot,j}}{\|\mathbf{X}_{\cdot,j}\|_2}$ if $\|\mathbf{X}_{\cdot,j}\|_2 \neq 0$, and setting the rest of the columns to zero.

- **Nuclear norm** is decomposable at every $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$. For a matrix \mathbf{X} with rank r and compact singular value decomposition $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^T$ where $\Sigma \in \mathbb{R}^{r \times r}$, we have $\mathbf{e} = \mathbf{U}\mathbf{V}^T$ and

$$\begin{aligned} T &= \{\mathbf{Y} \in \mathbb{R}^{n_1 \times n_2} : (\mathbf{I} - \mathbf{U}\mathbf{U}^T)\mathbf{Y}(\mathbf{I} - \mathbf{V}\mathbf{V}^T) = \mathbf{0}\} \\ &= \{\mathbf{Z}_1\mathbf{V}^T + \mathbf{U}\mathbf{Z}_2^T \mid \mathbf{Z}_1 \in \mathbb{R}^{n_1 \times r}, \mathbf{Z}_2 \in \mathbb{R}^{n_2 \times r}\}. \end{aligned}$$

The reader is referred to [36] for a more detailed discussion. Combining these examples with Table 1, it can be observed that the degrees of freedom for a sparse signal, column sparse matrix and a low rank matrix is same as the dimension of the supports for ℓ_1 norm, $\ell_{1,2}$ norm and nuclear norm respectively.




Figure 1: An example of a decomposable norm: ℓ_1 norm is decomposable at $\mathbf{x}_0 = (1, 0)$. The sign vector \mathbf{e} , the support T , and shifted subspace T^\perp are illustrated. A subgradient \mathbf{g} at \mathbf{x}_0 and its projection onto T^\perp are also shown.

Definition 2.2 Given a norm $\|\cdot\|$ decomposable at \mathbf{x} , define the constant κ as

$$\kappa = \left(\frac{\|\mathbf{e}\|_2}{L} \right)^2 \frac{n}{\dim(T)},$$

where \mathbf{e} and T are the sign and support of \mathbf{x} with respect to $\|\cdot\|$, and L is the Lipschitz constant of the norm, namely,

$$L = \sup_{\mathbf{z}_1 \neq \mathbf{z}_2 \in \mathbb{R}^n} \frac{\|\mathbf{z}_1\| - \|\mathbf{z}_2\|}{\|\mathbf{z}_1 - \mathbf{z}_2\|_2}. \quad (2.5)$$

The notation $\|\cdot\|_2$ denotes the Euclidean norm, i.e., the ℓ_2 norm for vectors and the Frobenius norm $\|\cdot\|_F$ for matrices. Note that L is a global property of the norm while \mathbf{e} , T and κ depend on both the norm and the point under consideration (decomposability is a local property in this sense). As it has been discussed in Lemma 2.2, the common norms ℓ_1 , $\ell_{1,2}$, and nuclear norm are decomposable at every point. The next lemma is an immediate consequence of Lemma 2.2 and states that κ is simply a bounded constant.

Lemma 2.3 *The value of κ for the ℓ_1 , $\ell_{1,2}$ and $\|\cdot\|_*$ norms are given as follows.*

- **ℓ_1 norm** has Lipschitz constant $L = \sup_{\mathbf{x} \neq 0} \frac{\|\mathbf{x}\|_1}{\|\mathbf{x}\|_2} = \sqrt{n}$. For a k -sparse vector \mathbf{x} , we have $\|\mathbf{e}\|_2 = \|\text{sgn}(\mathbf{x})\|_2 = \sqrt{k}$ and $\dim(T) = k$, which yield $\kappa = 1$ for any \mathbf{x} .
- **$\ell_{1,2}$ norm** has Lipschitz constant $L = \sup_{\mathbf{X} \neq 0} \frac{\|\mathbf{X}\|_{1,2}}{\|\mathbf{X}\|_F} = \sqrt{n_2}$. At any \mathbf{X} with k nonzero columns, we have $\|\mathbf{e}\|_F = \sqrt{k}$ and $\dim(T) = kn_1$ which gives $\kappa = 1$.
- **Nuclear norm** has a Lipschitz constant $L = \sup_{\mathbf{X} \neq 0} \frac{\|\mathbf{X}\|_*}{\|\mathbf{X}\|_F} = \sqrt{n_2}$ assuming $n_1 \geq n_2$. For a given matrix \mathbf{X} with rank r , we have $\dim(T) = r(n_1 + n_2 - r)$ and $\|\mathbf{e}\|_F = \sqrt{r}$. These give $\kappa = \frac{n_1}{n_1 + n_2 - r}$ which satisfies $\frac{1}{2} \leq \frac{n_1}{n_1 + n_2} \leq \kappa \leq 1$ for any \mathbf{X} .

2.1 Simultaneously structured models

We consider a signal \mathbf{x}_0 having several low dimensional structures S_1, S_2, \dots, S_τ simultaneously (e.g., sparsity, group sparsity, low-rank, etc.). Suppose these structures each correspond to a norm that promotes them (e.g., ℓ_1 , $\ell_{1,2}$, nuclear norm, etc.). We further assume these norms are decomposable (e.g., all the mentioned norms). While these assumptions may seem restrictive, they cover many cases of interest, for example all variations of the “sparse and low rank” matrices (see Section 3.2). We refer to such an \mathbf{x}_0 a *simultaneously structured model*.

To see why ℓ_1 norm is associated with sparsity, in Lemma 2.2, observe that the support T (for ℓ_1 norm) corresponds to the coordinates where \mathbf{x}_0 is nonzero and the dimension of the support is equal to the sparsity of the signal. Similarly, for the $\ell_{1,2}$ norm (the nuclear norm), dimension of the support depends only on the number of nonzero blocks (rank) of the matrix. In particular, the dimension of the support is equal to the degrees of freedom of a signal having the corresponding structure. In case of rank r matrices, this is equal to $r(n_1 + n_2 - r)$. Hence, all these norms are decomposable and are inherently connected to the corresponding low dimensional structure.

We will now introduce the relevant notation for a simultaneously structured signal \mathbf{x}_0 . Given $\{\|\cdot\|_{(i)}\}_{i=1}^\tau$ and \mathbf{x}_0 denote the corresponding supports by $\{T_i\}_{i=1}^\tau$ and the sign vectors by $\{\mathbf{e}_i\}_{i=1}^\tau$ (see Definition 2.1). Let $T_\cap = \bigcap_{i=1}^\tau T_i$ denote the joint support of \mathbf{x}_0 . Moreover, suppose \mathbf{x}_0 has other properties that can be expressed as a cone constraint $\mathbf{x}_0 \in \mathcal{C}$; an example is the positive semi-definiteness (PSD) constraint in the sparse phase retrieval problem mentioned in Section 1. Naturally, this additional information should help in recovery. To characterize the effect of the cone and the decomposable structures, consider the subspace

$$\mathcal{R} \triangleq T_\cap \cap \text{span}(\{\mathbf{y} \in \mathbb{R}^n \mid \langle \mathbf{x}_0, \mathbf{y} \rangle = 0, \mathbf{y} \in \mathcal{C}^*\})^\perp, \quad (2.6)$$

where $\text{span}(\cdot)$ returns the linear span of the elements of the set. Observe that $\mathbf{x}_0 \in \mathcal{R}$. When there is no cone constraint (i.e., $\mathcal{C} = \mathbb{R}^n$), we have $\mathcal{R} = T_\cap$. Similarly, if $\text{span}(\mathcal{C})$ is equal to \mathbb{R}^n and \mathbf{x}_0 lies in the interior of \mathcal{C} , we will again have $\mathcal{R} = T_\cap$. In general, the second term on the right hand side of (2.6) plays a critical role in our analysis when \mathbf{x}_0 lies on the boundary of \mathcal{C} . An example to this will be the “sparse and low rank matrices” which is discussed in Section 5. The following definition quantifies the angle of individual sign vectors with this subspace.

Definition 2.3 *Define $\mathbf{e}_{\cap,i} \triangleq \mathcal{P}_{\mathcal{R}}(\mathbf{e}_i)$ and let*

$$\theta_i \triangleq \frac{\|\mathbf{e}_{\cap,i}\|_2}{\|\mathbf{e}_i\|_2},$$

which is the cosine of the angle between \mathbf{e}_i and subspace \mathcal{R} .




Figure 2: A signal with two simultaneous structures and the associated definitions.

While κ_i captures the properties of the i th structure at the given point, θ_i reflects the relation between the different structures and the given cone. The quantities

$$\kappa = \min_{i=1,\dots,\tau} \kappa_i \quad , \quad \theta = \min_{i=1,\dots,\tau} \theta_i$$

will be used in the statement of results. Note that θ determines the maximum “spread” of vectors $\{\mathbf{e}_i\}_{i=1}^\tau$ from \mathcal{R} .

2.2 Convex recovery program

We will be investigating the recovery of the simultaneously structured \mathbf{x}_0 from its linear measurements $\mathcal{G}(\mathbf{x}_0)$. To recover the signal \mathbf{x}_0 , we would like to simultaneously minimize the norms $\|\cdot\|_{(i)}$, $i = 1, \dots, \tau$, which leads to a multi-objective (vector-valued) optimization problem. For all feasible points \mathbf{x} satisfying $\mathcal{G}(\mathbf{x}) = \mathcal{G}(\mathbf{x}_0)$ and $\mathbf{x} \in \mathcal{C}$, consider the set of achievable norms $\{\|\mathbf{x}\|_{(i)}\}_{i=1}^\tau$ denoted as points in \mathbb{R}^τ , as shown in Figure 3. Since the norms and the constraints are convex, the set of achievable values is also convex [38, Chapter 4]. The minimal points of this set, with respect to the partial order induced by the positive orthant \mathbb{R}_+^τ , form the *Pareto-optimal* front.

Definition 2.4 (Recoverability) *We call \mathbf{x}_0 recoverable if it is a Pareto optimal point; i.e., there does not exist any feasible $\mathbf{x}' \neq \mathbf{x}$ satisfying $\mathcal{G}(\mathbf{x}') = \mathcal{G}(\mathbf{x}_0)$ and $\mathbf{x}' \in \mathcal{C}$, with $\|\mathbf{x}'\|_{(i)} \leq \|\mathbf{x}_0\|_{(i)}$ for $i = 1, \dots, \tau$.*

Our vector-valued convex recovery program can be turned into a scalar optimization problem as

$$\begin{aligned} & \underset{\mathbf{x} \in \mathcal{C}}{\text{minimize}} && f(\mathbf{x}) = h(\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)}) \\ & \text{subject to} && \mathcal{G}(\mathbf{x}) = \mathcal{G}(\mathbf{x}_0), \end{aligned} \tag{2.7}$$

where $h : \mathbb{R}_+^\tau \rightarrow \mathbb{R}_+$ is chosen to be increasing with respect to the order induced by \mathbb{R}_+^τ . For convex problems with strong duality, it is known that the only scalarizing function h we need to produce all points \mathbf{x}_0 on the Pareto optimal front is the weighted sum $f(\mathbf{x}) = \sum_{i=1}^\tau \lambda_i \|\mathbf{x}\|_{(i)}$, where λ_i are positive weights and can be obtained as the coefficients of a supporting hyperplane at \mathbf{x}_0 (see, e.g., [38, Chapter 4]). Alternatively, another scalar objective function that can be used is the weighted maximum

$$f(\mathbf{x}) = \max_{i=1,\dots,\tau} \frac{1}{\|\mathbf{x}_0\|_{(i)}} \|\mathbf{x}\|_{(i)}. \tag{2.8}$$

We sometimes prefer to use this function instead of the weighted sum, since the weights corresponding to \mathbf{x}_0 are given explicitly. In particular, as it is discussed in Lemma 4.1, this function is the best choice to recover \mathbf{x}_0 via (2.7).

In Figure 3, consider the smallest m that makes \mathbf{x}_0 recoverable (i.e., whose corresponding achievable set has \mathbf{x}_0 as a Pareto optimal point). Then one can choose a function h and recover \mathbf{x}_0 by (2.7) using the m

measurements. If the number of measurements is any less, then *no* function can recover \mathbf{x}_0 . Our goal is to provide lower bounds on m .




Figure 3: Nested sets of achievable values shrink as the number of measurements grow, and they all contain \mathbf{x}_0 . We need at least m measurements for \mathbf{x}_0 to be recoverable, since for any $\underline{m} < m$, \mathbf{x}_0 is not on the Pareto optimal front.

Throughout this paper, we consider Gaussian measurements, defined as follows.

Definition 2.5 (Gaussian measurement operator) $\mathcal{G}(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a Gaussian measurement operator if $\mathcal{G}(\mathbf{x})$ is equivalent to the matrix multiplication $\mathbf{G}\mathbf{x}$ where $\mathbf{G} \in \mathbb{R}^{m \times n}$ has i.i.d. standard normal entries.

Note that in [8], Chandrasekaran et al. propose a general theory for constructing a suitable penalty, called an *atomic norm*, given a single set of atoms that describes the structure of the target object. In the case of simultaneous structures, this construction requires defining new atoms, and then ensuring the resulting atomic norm can be minimized in a computationally tractable way, which is nontrivial and often not easy. We briefly discuss such constructions as a future research direction in Section 7.

3 Main Results: Theorem Statements

In this section, we state our main theorems that aim to characterize the number of measurements needed to recover a simultaneously structured signal by convex or nonconvex programs. We first present our general results, followed by results for simultaneously sparse and low-rank matrices as a specific but important instance of the general case. The proofs are given in Section 4.

3.1 General simultaneously structured signals

This section deals with the recovery of a signal \mathbf{x}_0 that is simultaneously structured with S_1, S_2, \dots, S_τ as described in Section 2.1. We give a lower bound on the required number of measurements, that is determined by the support T_i with the smallest dimension. Before stating the theorem, we will give a relevant definition regarding the “size” of a cone.

Definition 3.1 (Cone width) Let \mathcal{M} be a closed convex cone in \mathbb{R}^n and let \mathbf{v} be a vector with i.i.d. standard normal entries. Then, the normalized width of \mathcal{M} is defined as

$$\eta(\mathcal{M}) = \frac{\mathbb{E}[\|\mathcal{P}_{\mathcal{M}}(\mathbf{v})\|_2]}{\sqrt{n}} \quad (3.1)$$

Note that we always have $\eta(\mathcal{M}) = \frac{\mathbb{E}[\|\mathcal{P}_{\mathcal{M}}(\mathbf{v})\|_2]}{\sqrt{n}} \leq \frac{\mathbb{E}[\|\mathbf{v}\|_2]}{\sqrt{n}} \leq 1$.

Theorem 3.1 Consider the programs in (2.7) using m generic measurements. Let

$$m_0 \triangleq \frac{n}{81\gamma^2} \inf_{\mathbf{g} \in \partial f(\mathbf{x}_0)} \frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2^2}{\|\mathbf{g}\|_2^2} - \tau$$

and $\gamma = \frac{2\eta(\mathcal{C})}{1-\eta(\mathcal{C}^\circ)}$. Then, whenever $m \leq \min\{\frac{n(1-\eta(\mathcal{C}^\circ))}{9}, m_0\}$, \mathbf{x}_0 will not be a minimizer for any of the programs in (2.7) with probability at least $1 - c_1 \exp(-c_2 \min\{m_0, n(1-\eta(\mathcal{C}^\circ))^2\})$ for some positive constants c_1 and c_2 .

In addition, observe that for a smaller cone \mathcal{C} , $\eta(\mathcal{C}^\circ)$ will be larger. Hence, it is reasonable to expect a smaller lower bound to the required number of measurements.

While this theorem is quite general, it is not very insightful as computing m_0 may not be easy. Next, we make an assumption that allows us to simplify the lower bound and give a general result for the programs in (2.7).

Assumption 1 For all $1 \leq i \neq j \leq \tau$, $\langle \mathbf{e}_{\cap,i}, \mathbf{e}_{\cap,j} \rangle \geq 0$ where $\{\mathbf{e}_{\cap,i}\}_{i=1}^\tau$ is given in Definition 2.3.

By Lemma 2.1, sign vectors $\{\mathbf{e}_i\}_{i=1}^\tau$ have positive inner products with \mathbf{x}_0 . Since, $\mathbf{x}_0 \in \mathcal{R}$, this also implies

$$\langle \mathbf{e}_{\cap,i}, \mathbf{x}_0 \rangle = \langle \mathbf{e}_i, \mathbf{x}_0 \rangle > 0 \quad \text{for } 1 \leq i \leq \tau. \quad (3.2)$$

Assumption 1 takes one step further and assumes pairwise nonnegative inner products between $\{\mathbf{e}_{\cap,i}\}_{i=1}^\tau$. The following lemma, which is easy to show, gives a sufficient condition for Assumption 1 to hold.

Lemma 3.1 Assumption 1 holds if the angle between \mathbf{x}_0 and \mathbf{e}_i is upper bounded by $\frac{\pi}{4}$ for all $1 \leq i \leq \tau$.

Angles between \mathbf{x}_0 and the \mathbf{e}_i 's are always upper bounded by $\frac{\pi}{2}$ due to Lemma 2.1, and when a stricter condition holds on these angles, Assumption 1 also holds.

As discussed before, there are various options for the scalarizing function in (2.7), with one choice being the weighted sum of norms. In fact, for a recoverable point \mathbf{x}_0 there always exists a weighted sum of norms which recovers it. This function is also often the choice in applications, where the space of positive weights is searched for a good combination. Thus, we can state the following theorem as a general result.

Theorem 3.2 Suppose Assumption 1 holds. Consider $d_{\min} = \min\{\dim(T_i) : i = 1, \dots, \tau\}$ and let γ, c_1, c_2 be same as in Theorem 3.1. Then, whenever $m \leq \min\{m'_0, \frac{n(1-\eta(\mathcal{C}^\circ))}{9}\}$, \mathbf{x}_0 will not be a minimizer of any of the recovery programs in (2.7) with probability at least $1 - c_1 \exp(-c_2 \min\{m'_0, n(1-\eta(\mathcal{C}^\circ))^2\})$, where

$$m'_0 \triangleq \frac{\kappa \theta^2}{81\gamma^2\tau} d_{\min} - \tau,$$

and

$$\kappa = \min_{1 \leq i \leq \tau} \frac{n\|\mathbf{e}_i\|_2^2}{L_i^2 \dim(T_i)} \quad \text{and} \quad \theta = \min_{1 \leq i \leq \tau} \frac{\|\mathbf{e}_{\cap,i}\|_2}{\|\mathbf{e}_i\|_2},$$

from Definitions 2.2 and 2.3 respectively.

Observe that Theorem 3.2 is stronger than stating “a particular function $h(\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)})$ will not work”. Instead, our result states that with high probability none of the programs in the class (2.7) can return \mathbf{x}_0 as the optimal unless the number of measurements are sufficiently large.

To understand the result better, note that the required number of measurements is proportional to d_{\min} which is the dimension of the smallest support. As we have argued in Section 2.1, dimension of the support corresponds to how structured the signal is. For sparse signals it is equal to the sparsity, and for a rank r matrix, it is equal to the degrees of freedom of the set of rank r matrices. Consequently, Theorem 3.2 suggests that even if the signal satisfies multiple structures, the required number of measurements is effectively determined by only one dominant structure.

Intuitively, the degrees of freedom of a simultaneously structured signal should be much lower, on the order of $\dim(T_\cap)$. Hence, there is a considerable gap between the expected measurements based on model complexity ($\dim(T_\cap)$) and the number of measurements needed for recovery via (2.7) ($\min_i \dim(T_i)$).

Finally, as shown in Section 5, κ and θ can be lower bounded by constants for the examples of norms we consider. For the specific cones considered there, $\eta(\mathcal{C}^\circ)$, $\eta(\mathcal{C})$ and γ are constants as well, (see Appendix A) and τ is generally a small positive integer. In these cases, the required number of measurements is directly determined by d_{\min} (see the next section).

3.2 Simultaneously Sparse and Low rank Matrices

We now focus on a special case, namely simultaneously sparse and low-rank (S&L) matrices. We consider matrices with nonzero entries contained in a small submatrix where the submatrix itself is low rank. Here, norms of interest are $\|\cdot\|_{1,2}$, $\|\cdot\|_1$ and $\|\cdot\|_\star$ and the cone of interest is the PSD cone. We also consider nonconvex approaches and contrast the results with convex approaches. For the nonconvex problem, we replace the norms $\|\cdot\|_1$, $\|\cdot\|_{1,2}$, $\|\cdot\|_\star$ with the functions $\|\cdot\|_0$, $\|\cdot\|_{0,2}$, $\text{rank}(\cdot)$ which gives the number of nonzero entries, the number of nonzero columns and rank of a matrix respectively and use the same cone constraint as the convex method. We show that convex methods perform poorly as predicted by the general result in Theorem 3.2, while nonconvex methods require optimal number of measurements (up to a logarithmic factor). Proofs are given in Section 5.

Definition 3.2 We say $\mathbf{X}_0 \in \mathbb{R}^{n_1 \times n_2}$ is an S&L matrix with (k_1, k_2, r) if the smallest submatrix that contains nonzero entries of \mathbf{X}_0 has size $k_1 \times k_2$ and $\text{rank}(\mathbf{X}_0) = r$. When \mathbf{X}_0 is symmetric, let $n = n_1 = n_2$ and $k = k_1 = k_2$. We consider the following cases.

- (a) General: $\mathbf{X}_0 \in \mathbb{R}^{n_1 \times n_2}$ is S&L with (k_1, k_2, r) .
- (b) PSD, arbitrary rank: $\mathbf{X}_0 \in \mathbb{R}^{n \times n}$ is PSD and S&L with (k, k, r) .
- (c) PSD, rank 1: $\mathbf{X}_0 = \mathbf{x}_0 \mathbf{x}_0^T$ where $\mathbf{x}_0 \in \mathbb{R}^n$ is k -sparse so that \mathbf{X}_0 is PSD and S&L with $(k, k, 1)$.

We are interested in S&L matrices with $k_1 \ll n_1, k_2 \ll n_2$ so that the matrix is sparse, and $r \ll \min\{k_1, k_2\}$ so that the submatrix containing the nonzero entries is low rank. Recall from Section 2.2 that our goal is to recover \mathbf{X}_0 from random Gaussian observations $\mathcal{G}(\mathbf{X}_0)$ via convex or nonconvex optimization programs. The measurements can be equivalently written as $\mathbf{G} \text{vec}(\mathbf{X}_0)$, where $\mathbf{G} \in \mathbb{R}^{m \times n_1 n_2}$ and $\text{vec}(\mathbf{X}_0) \in \mathbb{R}^{n_1 n_2}$ denotes the vector obtained by stacking the columns of \mathbf{X}_0 .

Based on the results in Section 3.1, we obtain lower bounds on the number of measurements for convex recovery. We additionally show that significantly fewer measurements are sufficient for non-convex programs to uniquely recover \mathbf{X}_0 ; thus proving a performance gap between convex and nonconvex approaches. The following theorem summarizes the results.

Theorem 3.3 (Performance of S&L matrix recovery) Assume $m \leq \frac{n_1 n_2}{9}$, and consider recovering $\mathbf{X}_0 \in \mathbb{R}^{n_1 \times n_2}$ via

$$\underset{\mathbf{X} \in \mathcal{C}}{\text{minimize}} f(\mathbf{X}) \quad \text{subject to} \quad \mathcal{G}(\mathbf{X}) = \mathcal{G}(\mathbf{X}_0). \quad (3.3)$$

For the cases given in Definition 3.2, the following convex and nonconvex recovery results hold for some positive constants c_1, c_2 .

- (a) General model:
 - (a1) Let $f(\mathbf{X}) = \|\mathbf{X}\|_{1,2} + \lambda_1 \|\mathbf{X}^T\|_{1,2} + \lambda_2 \|\mathbf{X}\|_\star$ and $\mathcal{C} = \mathbb{R}^{n_1 \times n_2}$. Then, (3.3) will fail to recover \mathbf{X}_0 with probability $1 - \exp(-c_1 d)$ for all possible $\lambda_1, \lambda_2 \geq 0$ whenever $m \leq c_2 d$ where $d = \min\{n_1 k_2, n_2 k_1, (n_1 + n_2)r\}$.
 - (a2) Let $f(\mathbf{X}) = \frac{\|\mathbf{X}\|_{0,2}}{k_2} + \frac{\|\mathbf{X}^T\|_{0,2}}{k_1} + \frac{\text{rank}(\mathbf{X})}{r}$ and $\mathcal{C} = \mathbb{R}^{n_1 \times n_2}$. Then, (3.3) will uniquely recover \mathbf{X}_0 with probability $1 - \exp(-c_1 m)$ whenever $m \geq c_2 \max\{(k_1 + k_2)r, k_1 \log \frac{n_1}{k_1}, k_2 \log \frac{n_2}{k_2}\}$.

Setting	Nonconvex sufficient m	Convex required m
General model	$O(\max\{rk, k \log \frac{n}{k}\})$	$\Omega(rn)$
PSD, arbitrary rank	$O(\max\{rk, k \log \frac{n}{k}\})$	$\Omega(rn)$
PSD, rank 1	$O(k \log \frac{n}{k})$	$\Omega(\min\{k^2, n\})$

Table 2: Summary of recovery results for models in Definition 3.2, assuming $n_1 = n_2 = n$ and $k_1 = k_2 = k$. For the PSD rank 1 case, we assume $\frac{\|\mathbf{x}_0\|_1}{\sqrt{k}\|\mathbf{x}_0\|_2}$ to be a constant. Nonconvex approaches are optimal up to a logarithmic factor, while convex approaches perform poorly.

(b) *PSD, arbitrary rank:*

(b1) Let $f(\mathbf{X}) = \|\mathbf{X}\|_{1,2} + \lambda\|\mathbf{X}\|_*$ and $\mathcal{C} = \mathbb{S}_+^n$. Then, (3.3) will fail to recover \mathbf{X}_0 with probability $1 - \exp(-c_1 rn)$ for all possible $\lambda \geq 0$ whenever $m \leq c_2 rn$.

(b2) Let $f(\mathbf{X}) = \frac{2\|\mathbf{X}\|_{0,2}}{k} + \frac{\text{rank}(\mathbf{X})}{r}$ and $\mathcal{C} = \mathbb{S}^n$. Then, (3.3) will uniquely recover \mathbf{X}_0 with probability $1 - \exp(-c_1 m)$ whenever $m \geq c_2 \max\{rk, k \log \frac{n}{k}\}$.

(c) *PSD, rank 1:*

(c1) Let $f(\mathbf{X}) = \|\mathbf{X}\|_1 + \lambda\|\mathbf{X}\|_*$ and $\mathcal{C} = \mathbb{S}_+^n$. Then, (3.3) will fail to recover \mathbf{X}_0 with probability $1 - \exp(-c_1 d)$ for all possible $\lambda \geq 0$ whenever $m \leq c_2 d$ where $d = \frac{\|\mathbf{x}_0\|_1^2}{k\|\mathbf{x}_0\|_2} \min\{k^2, n\}$.

(c2) Let $f(\mathbf{X}) = \frac{\|\mathbf{X}\|_0}{k^2} + \frac{\text{rank}(\mathbf{X})}{r}$ and $\mathcal{C} = \mathbb{S}^n$. Then, (3.3) will uniquely recover \mathbf{X}_0 with probability $1 - \exp(-c_1 m)$ whenever $m \geq c_2 k \log \frac{n}{k}$.

The nonconvex programs require almost the same number of measurements as the degrees of freedom (or number of parameters) of the underlying model. For instance, it is known that the degrees of freedom of a rank r matrix of size $k_1 \times k_2$ is simply $r(k_1 + k_2 - r)$ which is $O((k_1 + k_2)r)$. Hence, the nonconvex results are optimal up to a logarithmic factor. On the other hand, our results on the convex programs that follow from Theorem 3.2 indicate that the required number of measurements are determined by the support with the smallest dimension. For example, for \mathbf{X}_0 obeying the ‘general model’, the dimension of supports for the norms $\|\mathbf{X}\|_{1,2}$, $\|\mathbf{X}^T\|_{1,2}$ and $\|\mathbf{X}\|_*$ at \mathbf{X}_0 are $n_1 k_2$, $n_2 k_1$ and $(n_1 + n_2 - r)r$ respectively, and we indeed require $\Omega(\min\{n_1 k_2, n_2 k_1, (n_1 + n_2)r\})$ measurements. Table 2 provides a quick comparison of the results on S&L. It can be seen that there is a meaningful gap: while nonconvex programs are successful with orderwise optimal number of measurements, convex programs need significantly more measurements. We observe a similar gap for the special case of simultaneously sparse and rank-1 PSD matrices.

While there is a gap between nonconvex and convex recovery, our performance bound on the mixed convex optimization, can be related to the performance of best of the individual norms. This comes from the fact that for many structure inducing norms, dimension of the support and the required number of measurements for recovery are same up to a logarithmic factor. For instance, for the “PSD, rank 1” model, only nuclear norm minimization would require $O(n)$ and only ℓ_1 norm would require $O(k^2 \log \frac{n}{k})$ measurements, where the dimensions of the corresponding supports are of order n and k^2 respectively. Hence, the best of the individual norms gives $O(\min\{n, k^2 \log \frac{n}{k}\})$, which is only logarithmically larger than the bound given by Theorem 3.3, statement (c1). Similar result is true for the matrices obeying the “PSD, arbitrary rank” model, when we consider the fact that minimizing only $\ell_{1,2}$ norm would require $O(kn)$ ([36]) and only nuclear norm would require $O(rn)$ measurements for perfect recovery. Hence the best individual norm requires $O(\min\{kn, rn\}) = O(rn)$, which is same as the bound given in statement (b1). Overall, mixed norms are not able to significantly outperform the best of the individual norms.

Finally, as we saw in Section 3.1, adding a cone constraint to the recovery program does not help in reducing the lower bound by more than a constant factor. In particular, we discuss the positive semidefiniteness assumption, and show that in the sparse phase retrieval problem, the number of measurements remain high even when we include this extra information. On the other hand, the nonconvex recovery programs performs well even without the PSD constraint.

4 General Simultaneously Structured Model Recovery

Recall the setup from Section 2. As discussed in Section 2.1, we consider a vector $\mathbf{x}_0 \in \mathbb{R}^n$ at which a family of norms $\{\|\cdot\|_{(i)}\}_{i=1}^\tau$ are decomposable, and \mathbf{x}_0 satisfies the cone constraint $\mathbf{x}_0 \in \mathcal{C}$. Recall that the supports corresponding to the norms are given by $\{T_i\}_{i=1}^\tau$ and $T_\cap = \bigcap_{i=1}^\tau T_i$. This section is dedicated to the proofs of theorems in Section 3.1 and additional side results where the goal is to find lower bounds on the required number of measurements to recover \mathbf{x}_0 . By discussions in Section 2.2, we can focus on point-wise weighted summation of norms and the same lower bounds will work for any function in (2.7).

4.1 Preliminary Lemmas

We first show that, to recover \mathbf{x}_0 , the objective function $\max_{1 \leq i \leq \tau} \frac{\|\mathbf{x}\|_{(i)}}{\|\mathbf{x}_0\|_{(i)}}$ can be viewed as the ‘best’ among the functions mentioned in (2.7).

Lemma 4.1 *Consider the class of recovery programs in (2.7). If the program*

$$\begin{aligned} & \underset{\mathbf{x} \in \mathcal{C}}{\text{minimize}} && f_{\text{best}}(\mathbf{x}) \triangleq \max_{i=1, \dots, \tau} \frac{\|\mathbf{x}\|_{(i)}}{\|\mathbf{x}_0\|_{(i)}} \\ & \text{subject to} && \mathcal{G}(\mathbf{x}) = \mathcal{G}(\mathbf{x}_0) \end{aligned} \quad (4.1)$$

fails to recover \mathbf{x}_0 , then any member of this class will also fail to recover \mathbf{x}_0 .

Proof. Suppose (4.1) does not have \mathbf{x}_0 as an optimal solution and there exists \mathbf{x}' such that $f_{\text{best}}(\mathbf{x}') \leq f_{\text{best}}(\mathbf{x}_0)$, then

$$\frac{1}{\|\mathbf{x}_0\|_{(i)}} \|\mathbf{x}'\|_{(i)} \leq f_{\text{best}}(\mathbf{x}') \leq f_{\text{best}}(\mathbf{x}_0) = 1, \quad \text{for } i = 1, \dots, \tau,$$

which implies,

$$\|\mathbf{x}'\|_{(i)} \leq \|\mathbf{x}_0\|_{(i)}, \quad \text{for all } i = 1, \dots, \tau. \quad (4.2)$$

Conversely, given (4.2), we have $f_{\text{best}}(\mathbf{x}') \leq f_{\text{best}}(\mathbf{x}_0)$ from the definition of f_{best} .

Furthermore, since we assume $h(\cdot)$ in (2.7) is non-decreasing in its arguments and increasing in at least one of them, (4.2) implies $f(\mathbf{x}') \leq f(\mathbf{x}_0)$ for any such function $f(\cdot)$. Thus, failure of $f_{\text{best}}(\cdot)$ in recovery of \mathbf{x}_0 implies failure of any other function in (2.7) in this task. ■

The following lemma gives necessary conditions for \mathbf{x}_0 to be a minimizer of the problem (2.7).

Lemma 4.2 *Let \mathcal{G}^* denote the adjoint of the linear map \mathcal{G} . If \mathbf{x}_0 is a minimizer of the program (2.7), then there exist $\mathbf{v} \in \mathcal{C}^*$, \mathbf{z} , and $\mathbf{g} \in \partial f(\mathbf{x}_0)$ such that*

$$\mathbf{g} - \mathbf{v} - \mathcal{G}^*(\mathbf{z}) = 0 \quad \text{and} \quad \langle \mathbf{x}_0, \mathbf{v} \rangle = 0.$$

The proof of Lemma 4.2 follows from the KKT conditions for (2.7) to have \mathbf{x}_0 as an optimal solution [39, Section 4.7].

The next lemma describes the subdifferential of any general function $f(\mathbf{x}) = h(\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)})$ as discussed in Section 2.2.

Lemma 4.3 *For any subgradient of the function $f(\mathbf{x}) = h(\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)})$ at $\mathbf{x} \neq 0$ defined by convex function $h(\cdot)$, there exists non-negative constants w_i , $i = 1, \dots, \tau$ such that*

$$\mathbf{g} = \sum_{i=1}^{\tau} w_i \mathbf{g}_i$$

where $\mathbf{g}_i \in \partial \|\mathbf{x}_0\|_{(i)}$.

Proof. Consider the function $N(\mathbf{x}) = [\|\mathbf{x}\|_{(1)}, \dots, \|\mathbf{x}\|_{(\tau)}]^T$ by which we have $f(\mathbf{x}) = h(N(\mathbf{x}))$. By Theorem 10.49 in [40] we have

$$\partial f(\mathbf{x}) = \bigcup \{ \partial(\mathbf{y}^T N(\mathbf{x})) : \mathbf{y} \in \partial h(N(\mathbf{x})) \}$$

where we used the convexity of f and h . Now notice that any $\mathbf{y} \in \partial h(N(\mathbf{x}))$ is a non-negative vector because of the monotonicity assumption on $h(\cdot)$. This implies that any subgradient $\mathbf{g} \in \partial f(\mathbf{x})$ is in the form of $\partial(\mathbf{w}^T N(\mathbf{x}))$ for some nonnegative vector \mathbf{w} . The desired result simply follows because subgradients of conic combination of norms are conic combinations of their subgradients, (see e.g. [41]). ■

Using Lemmas 4.2 and 4.3, we now provide the proofs of Theorems 3.1 and 3.2.

4.2 Proof of Theorem 3.1

Suppose \mathbf{x}_0 is a minimizer of (2.7). From Lemma 4.2, there exist a $\mathbf{g} \in \partial f(\mathbf{x}_0)$, $\mathbf{z} \in \mathbb{R}^m$ and $\mathbf{v} \in \mathcal{C}^*$ such that

$$\mathbf{g} = \mathcal{G}^*(\mathbf{z}) + \mathbf{v} \quad (4.3)$$

and $\langle \mathbf{x}_0, \mathbf{v} \rangle = 0$. To use the spectral properties of random Gaussian map \mathcal{G} we will eliminate the contribution of \mathbf{v} in equation (4.3). Recalling, (2.6), observe that $\mathcal{P}_{\mathcal{R}}(\mathbf{v}) = 0$ as $\mathbf{v} \in \mathcal{R}^\perp$ by definition. Now, projecting both sides of (4.3) onto the subspace \mathcal{R} gives

$$\mathcal{P}_{\mathcal{R}}(\mathbf{g}) = \mathcal{P}_{\mathcal{R}}(\mathcal{G}^*(\mathbf{z})). \quad (4.4)$$

From Lemma 4.3, $\mathcal{P}_{\mathcal{R}}(\mathbf{g})$ lies in the span of $\{\mathbf{e}_{\cap, i}\}_{i=1}^\tau$ which is a τ dimensional subspace. On the other hand $\text{range}(\mathcal{P}_{\mathcal{R}}(\mathcal{G}^*))$ is an m dimensional subspace chosen uniformly at random in \mathcal{R} . Hence, whenever $m \leq \dim(\mathcal{R}) - \tau$, these subspaces have trivial intersection with probability 1, which implies that there does not exist a \mathbf{z} satisfying (4.4). Hence, for recovery of \mathbf{x}_0 , we need $m > \dim(\mathcal{R}) - \tau$.

Let us call $m' = \min\{m_0, \frac{n(1-\eta(\mathcal{C}^\circ))}{9}\}$. If $\dim(\mathcal{R}) - \tau \geq m'$, we can already conclude with the desired result. Otherwise, we assume $m' \geq \dim(\mathcal{R}) - \tau$. Furthermore, it is safe to prove the result for $m = m'$, as fewer measurements can only increase the chance of failure. Next, we consider three events each of which hold with high probability. Below, c'_1, c'_2 are the proper corresponding constants for each case.

- Using Theorem A.3, since $m \leq \frac{7(1-\eta(\mathcal{C}^\circ))n}{16}$, with probability at least $1 - c'_1 \exp(-c'_2(1 - \eta(\mathcal{C}^\circ))^2 n)$, for all $\mathbf{z} \in \mathbb{R}^m$, we have

$$\|\mathcal{G}^*(\mathbf{z})\|_2 \leq \gamma \|\mathcal{P}_{\mathcal{C}}(\mathcal{G}^*(\mathbf{z}))\|_2 \quad (4.5)$$

- Observe that \mathcal{G}^* is equivalent to an $n \times m$ matrix with i.i.d. standard normal entries and similarly $\mathcal{P}_{\mathcal{R}}(\mathcal{G}^*)$ is equivalent to an $\dim(\mathcal{R}) \times m$ matrix with i.i.d. standard normal entries after proper unitary transformation. Using Corollary (5.35) in [50] and choosing $t = \sqrt{m'}$, with probability $1 - 4 \exp(-\frac{m'}{2})$, we have

$$\sigma_{\min}(\mathcal{G}^*) > \sqrt{n} - \sqrt{m'} - \sqrt{m'} \geq \frac{\sqrt{n}}{3} \quad (4.6)$$

$$\sigma_{\max}(\mathcal{P}_{\mathcal{R}}(\mathcal{G}^*)) < \sqrt{\dim(\mathcal{R})} + \sqrt{m'} + \sqrt{m'} \leq 3\sqrt{m_0 + \tau} \quad (4.7)$$

Overall, inequalities (4.5), (4.6) and (4.7) hold with probability at least $1 - c_1 \exp(-c_2 \min\{m_0, n(1 - \eta(\mathcal{C}^\circ))^2\})$. Assuming they hold, we will show that multipliers $\mathbf{g}, \mathbf{v}, \mathbf{z}$ in (4.3) cannot exist. To show this by contradiction, assume these inequalities hold and such $\mathbf{g}, \mathbf{v}, \mathbf{z}$ exist.

Since $\mathbf{v} \in \mathcal{C}^*$, from Lemma A.1 we have $\mathcal{P}_{\mathcal{C}}(-\mathbf{v}) = \mathcal{P}_{\mathcal{C}}(\mathcal{G}^*(\mathbf{z}) - \mathbf{g}) = 0$. Using Corollary A.1,

$$\|\mathbf{g}\|_2 \geq \|\mathcal{P}_{\mathcal{C}}(\mathcal{G}^*(\mathbf{z}))\|_2. \quad (4.8)$$

Now, combining (4.4), (4.5) and (4.8), we have

$$\frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2}{\sigma_{\max}(\mathcal{P}_{\mathcal{R}}(\mathcal{G}^*))} \leq \|\mathbf{z}\|_2 \leq \frac{\gamma\|\mathbf{g}\|_2}{\sigma_{\min}(\mathcal{G}^*)}. \quad (4.9)$$

Using (4.6) and (4.7), we can rewrite (4.9) as

$$\frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2}{3\sqrt{m_0} + \tau} < \|\mathbf{z}\|_2 < \frac{3\gamma\|\mathbf{g}\|_2}{\sqrt{n}},$$

which gives

$$m_0 > n \left(\frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2}{9\gamma\|\mathbf{g}\|_2} \right)^2 - \tau \geq \inf_{\mathbf{g} \in \partial f(\mathbf{x}_0)} n \left(\frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2}{9\gamma\|\mathbf{g}\|_2} \right)^2 - \tau = m_0,$$

which is a contradiction.

4.3 Proof of Theorem 3.2

The result follows immediately from Theorem 3.1. We first find a lower bound on m_0 . From Lemma 4.3, any $\mathbf{g} \in f(\mathbf{x}_0)$ can be written as $\mathbf{g} = \sum_{i=1}^{\tau} w_i \mathbf{g}_i$ for some non-negative coefficients w_i . Now, using Lemma B.1, we can bound the subgradient as

$$\|\mathbf{g}\|_2 = \left\| \sum_{i=1}^{\tau} w_i \mathbf{g}_i \right\|_2 \leq \sum_{i=1}^{\tau} w_i \|\mathbf{g}_i\|_2 \leq \sum_{i=1}^{\tau} w_i L_i,$$

where L_i is the Lipschitz constant of norm $\|\cdot\|_{(i)}$. Next, from the definition of κ we have

$$\frac{1}{n} \left(\sum_{i=1}^{\tau} w_i L_i \right)^2 \leq \frac{\tau}{n} \sum_{i=1}^{\tau} w_i^2 L_i^2 = \frac{\tau}{n} \sum_{i=1}^{\tau} w_i^2 \frac{n\|\mathbf{e}_i\|_2^2}{\kappa_i \dim(T_i)} \leq \frac{\tau}{\kappa d_{\min}} \sum_{i=1}^{\tau} w_i^2 \|\mathbf{e}_i\|_2^2,$$

where the leftmost inequality follows from the Cauchy-Schwarz inequality. Overall, we find

$$\frac{n}{\|\mathbf{g}\|_2^2} \geq \frac{\kappa d_{\min}}{\tau \sum_{i=1}^{\tau} w_i^2 \|\mathbf{e}_i\|_2^2}. \quad (4.10)$$

We will now estimate $\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2$. First, we have

$$\mathcal{P}_{\mathcal{R}}(\mathbf{g}) = \sum_{i=1}^{\tau} w_i \mathcal{P}_{\mathcal{R}}(\mathbf{g}_i) = \sum_{i=1}^{\tau} w_i \mathbf{e}_{\cap, i},$$

which by using Assumption 1 gives

$$\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2^2 = \left\| \sum_{i=1}^{\tau} w_i \mathbf{e}_{\cap, i} \right\|_2^2 \geq \sum_{i=1}^{\tau} w_i^2 \|\mathbf{e}_{\cap, i}\|_2^2 \geq \theta^2 \sum_{i=1}^{\tau} w_i^2 \|\mathbf{e}_i\|_2^2. \quad (4.11)$$

Combining (4.10) and (4.11), we find

$$n \left(\frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2}{9\gamma\|\mathbf{g}\|_2} \right)^2 \geq \frac{\kappa \theta^2}{81\gamma^2 \tau} d_{\min} = m'_0 + \tau. \quad (4.12)$$

The last inequality is true for all $\mathbf{g} \in \partial f(\mathbf{x}_0)$; hence,

$$m_0 + \tau = \min_{\mathbf{g} \in \partial f(\mathbf{x}_0)} n \left(\frac{\|\mathcal{P}_{\mathcal{R}}(\mathbf{g})\|_2}{9\gamma\|\mathbf{g}\|_2} \right)^2 \geq m'_0 + \tau.$$

Theorem 3.2 follows immediately from Theorem 3.1 as any $m \leq \min\{m'_0, \frac{n(1-\eta(\mathcal{C}^\circ))}{9}\}$ will satisfy $m \leq \min\{m_0, \frac{n(1-\eta(\mathcal{C}^\circ))}{9}\}$ as well.

5 Simultaneously Sparse and Low-rank Matrices

Using the general framework provided in Section 3.1, in this section we present the proof of Theorem 3.3, which states various convex and nonconvex recovery results for the S&L models. We divide the proof into two parts, namely proofs for convex recovery and nonconvex recovery. We begin with convex recovery results.

5.1 Convex recovery results for S&L

In this section, we prove the statements of Theorem 3.3 regarding convex approaches, using Theorem 3.2. We begin with the following lemma which gives results on sign vectors and supports for the S&L model. The proof is provided in Appendix B.

Lemma 5.1 *Denote the norm $f(\mathbf{X}) = \|\mathbf{X}^T\|_{1,2}$ for all $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ by $\|\cdot\|_{1,2}^T$. Given a matrix $\mathbf{X}_0 \in \mathbb{R}^{n_1 \times n_2}$, let $\mathbf{E}_\star, \mathbf{E}_c, \mathbf{E}_r, \mathbf{E}_1$ and T_\star, T_c, T_r, T_1 be the sign vectors and supports for the norms $\|\cdot\|_\star, \|\cdot\|_{1,2}, \|\cdot\|_{1,2}^T, \|\cdot\|_1$ respectively. Then,*

- $\mathbf{E}_\star, \mathbf{E}_r, \mathbf{E}_c \in T_\star \cap T_c \cap T_r$,
- $\langle \mathbf{E}_\star, \mathbf{E}_r \rangle \geq 0, \langle \mathbf{E}_\star, \mathbf{E}_c \rangle \geq 0$, and $\langle \mathbf{E}_c, \mathbf{E}_r \rangle \geq 0$.

Now, assume $\mathbf{X}_0 = \sigma \mathbf{u} \mathbf{v}^T$ is a rank one matrix, where $\mathbf{u} \in \mathbb{R}^{n_1}$ and $\mathbf{v} \in \mathbb{R}^{n_2}$ are k_1 and k_2 sparse unit length vectors and $\sigma = \|\mathbf{X}_0\|_F$. We have

- $\mathbf{E}_\star \in T_\star \cap T_1$ and $\|\mathcal{P}_{T_\star \cap T_1}(\mathbf{E}_1)\|_F \geq \max\{\frac{\|\mathbf{u}\|_1}{\sqrt{k_1}}, \frac{\|\mathbf{v}\|_1}{\sqrt{k_2}}\} \|\mathbf{E}_1\|_F$,
- $\langle \mathcal{P}_{T_\star \cap T_1}(\mathbf{E}_1), \mathbf{E}_\star \rangle \geq 0$.

5.1.1 Proof of Theorem 3.3: Convex cases

Proof of (a1) We use the functions $\|\cdot\|_{1,2}, \|\cdot\|_{1,2}^T$ and $\|\cdot\|_\star$ without the cone constraint, i.e., $\mathcal{C} = \mathbb{R}^{n_1 \times n_2}$. Following the notation of Lemma 5.1, $T_\cap = T_\star \cap T_c \cap T_r$ and all the sign vectors lie on $\mathcal{R} = T_\cap$, which means $\theta = 1$, and they have pairwise nonnegative inner products. Also, $\dim(T_c) = k_2 n_1$, $\dim(T_r) = k_1 n_2$ and $\dim(T_\star) = (n_1 + n_2 - r)r$. Hence $d_{\min} = \min\{(n_1 + n_2 - r)r, n_1 k_2, n_2 k_1\}$. Furthermore, from Lemma 2.3, $\kappa \geq \frac{1}{2}$. Applying Theorem 3.2 gives the result.

Proof of (b1) In this case, we apply Lemma B.3. We have $\mathcal{R} = T_\cap \cap \mathbb{S}^n$, the norms are the same as in the general model, and $\theta \geq \frac{1}{\sqrt{2}}$. Also, pairwise inner products are positive, $d_{\min} = \min\{(2n - r)r, kn\}$ and $\kappa \geq \frac{1}{2}$. Based on Corollary A.2, for the PSD cone we have $\gamma \leq 11$. The result follows from Theorem 3.2.

Proof of (c1) In this case, the norms are $\|\cdot\|_1$ and $\|\cdot\|_\star$. From Lemma B.3, $\mathcal{R} = T_1 \cap T_\star \cap \mathbb{S}^n$ and $\theta \geq \frac{\|\mathbf{x}_0\|_1}{\|\mathbf{x}_0\|_2 \sqrt{k}}$. Similar to (b1), $\gamma \leq 11$. Also, $d_{\min} = \min\{k^2, n\}$. The result follows from Theorem 3.2.

5.2 Nonconvex recovery results for S&L

We first state a lemma that will be useful in proving the nonconvex results. The proof is provided in the Appendix C and uses standard arguments.

Lemma 5.2 *Consider the set of matrices S in $\mathbb{R}^{n_1 \times n_2}$ that are supported over a $d_1 \times d_2$ submatrix with rank at most s . There exists a constant $c > 0$ such that whenever $m \geq c \min\{(d_1 + d_2)s, d_1 \log \frac{n_1}{d_1}, d_2 \log \frac{n_2}{d_2}\}$, with probability $1 - 2 \exp(-cm)$ an i.i.d. Gaussian operator will satisfy the following,*

$$\mathcal{G}(\mathbf{X}) \neq 0, \quad \text{for all } \mathbf{X} \in S. \quad (5.1)$$

5.2.1 Proof of Theorem 3.3: Nonconvex cases

Denote the sphere in $\mathbb{R}^{n_1 \times n_2}$ with unit Frobenius norm by \mathcal{B} .

Proof of (a2) Observe that the function $f(\mathbf{X}) = \frac{\|\mathbf{X}\|_{0,2}}{\|\mathbf{X}_0\|_{0,2}} + \frac{\|\mathbf{X}^T\|_{0,2}}{\|\mathbf{X}_0^T\|_{0,2}} + \frac{\text{rank}(\mathbf{X})}{\text{rank}(\mathbf{X}_0)}$ satisfies the triangle inequality and we have $f(\mathbf{X}_0) = 3$. Hence, if all null space elements $\mathbf{W} \in \mathcal{N}(\mathcal{G})$ satisfy $f(\mathbf{W}) > 6$, we have

$$f(\mathbf{X}) \geq f(\mathbf{X} - \mathbf{X}_0) - f(-\mathbf{X}_0) > 3,$$

for all feasible \mathbf{X} which implies \mathbf{X}_0 being the unique minimizer.

Consider the set S of matrices, which are supported over a $6k_1 \times 6k_2$ submatrix with rank at most $6r$. Observe that any \mathbf{Z} satisfying $f(\mathbf{Z}) \leq 6$ belongs to S . Hence ensuring $\mathcal{N}(\mathcal{G}) \cap S = \{0\}$ would ensure $f(\mathbf{W}) > 6$ for all $\mathbf{W} \in \mathcal{N}(\mathcal{G})$. Since S is a cone, this is equivalent to $\mathcal{N}(\mathcal{G}) \cap (S \cap \mathcal{B}) = \emptyset$. Now, applying Lemma 5.2 with set S and $d_1 = 6k_1$, $d_2 = 6k_2$, $s = 6r$ we find the desired result.

Proof of (b2) Observe that due to the symmetry constraint,

$$f(\mathbf{X}) = \frac{\|\mathbf{X}\|_{0,2}}{\|\mathbf{X}_0\|_{0,2}} + \frac{\|\mathbf{X}^T\|_{0,2}}{\|\mathbf{X}_0^T\|_{0,2}} + \frac{\text{rank}(\mathbf{X})}{\text{rank}(\mathbf{X}_0)}.$$

Hence, the minimization is the same as (a2), the matrix is rank r contained in a $k \times k$ submatrix and we additionally have the positive semidefinite constraint which can only reduce the amount of required measurements compared to (a2). Consequently, the result follows by applying Lemma 5.2, similar to (a2).

Proof of (c2) Let $C = \{\mathbf{X} \neq 0 \mid f(\mathbf{X}) \leq f(\mathbf{X}_0)\}$. Since $\text{rank}(\mathbf{X}_0) = 1$, if $f(\mathbf{X}) \leq f(\mathbf{X}_0) = 2$, $\text{rank}(\mathbf{X}) = 1$. With the symmetry constraint, this means $\mathbf{X} = \pm \mathbf{x}\mathbf{x}^T$ for some l -sparse \mathbf{x} . Observe that $\mathbf{X} - \mathbf{X}_0$ has rank at most 2 and is contained in a $2k \times 2k$ submatrix as $l \leq k$. Let S be the set of matrices that are symmetric and whose support lies in a $2k \times 2k$ submatrix. Using Lemma 5.2 with $s = 2$, $d_1 = d_2 = 2k$, whenever $m \geq ck \log \frac{n}{k}$, with desired probability all nonzero $\mathbf{W} \in S$ will satisfy $\mathcal{A}(\mathbf{W}) \neq 0$. Consequently, any $\mathbf{X} \in C$ will have $\mathcal{A}(\mathbf{X}) \neq \mathcal{A}(\mathbf{X}_0)$, hence \mathbf{X}_0 will be the unique minimizer.

6 Numerical Experiments

In this section, we numerically verify our theoretical bounds on the number of measurements for the Sparse and Low-rank recovery problem. We demonstrate the empirical performance of the weighted maximum of the norms f_{best} (see Lemma 4.1), as well as the weighted sum of norms.

The experimental setup is as follows. Our goal is to explore how the number of required measurements m (or m/θ^2 in the second set of experiments) scales with the size of the matrix n . We consider a grid of (m, n) values, and generate at least 100 test instances for each grid point (in the boundary areas, we increase the number of instances to at least 200). We generate the target matrix \mathbf{X}_0 by generating a $k \times r$ i.i.d. Gaussian matrix \mathbf{G} , and inserting the $k \times k$ matrix $\mathbf{G}\mathbf{G}^T$ in an $n \times n$ matrix of zeros. We take $r = 1$ and $k = 8$ in all of the following experiments; even with these small values, we can observe the scaling predicted by our bounds. In each test, we measure the normalized recovery error $\frac{\|\mathbf{X} - \mathbf{X}_0\|_F}{\|\mathbf{X}_0\|_F}$ and declare successful recovery when this error is less than 10^{-4} . The optimization programs are solved using the CVX package [42], which calls the SDP solver SeDuMi [43].

We first test our bound in part (b) of Theorem 3.3, $\Omega(nr)$, on the number of measurements for recovery in the case of minimizing $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_{1,2}}{\|\mathbf{X}_0\|_{1,2}}\}$ over the set of positive semi-definite matrices. Figure 4 shows the results, which demonstrates m scaling linearly with n (note that $r = 1$).

Next, we replace $\ell_{1,2}$ norm with ℓ_1 norm and consider a recovery program that emphasizes entry-wise sparsity rather than block sparsity. Figure 5 demonstrates the lower bound $\Omega(\min\{k^2, n\})$ in Part (c) of Theorem 3.3 where we attempt to recover a rank-1 positive semi-definite matrix \mathbf{X}_0 by minimizing $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$




Figure 4: Performance of the recovery program minimizing $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_{1,2}}{\|\mathbf{X}_0\|_{1,2}}\}$ with a PSD constraint. The dark region corresponds to the experimental region of failure due to insufficient measurements. As predicted by Theorem 3.3, the number of required measurements increases linearly with nr .

subject to the measurements and a PSD constraint. As pointed out in Section 5, the lower bound given in Theorem 3.2 has a target-dependent term θ^2 , so to make the different test instances comparable, we plot $\frac{m}{\theta^2}$ versus n . The green curve in the figure shows the empirical 95% failure boundary, depicting the region of failure with high probability that our results have predicted. It starts off growing linearly with n , when the term nr dominates the term k^2 , and then saturates as n grows and the k^2 term (which is a constant in our experiments) becomes dominant.




Figure 5: Performance of the recovery program minimizing $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ with a PSD constraint. $r = 1, k = 8$ and n is allowed to vary. The plot shows $\frac{m}{\theta^2}$ versus n to better illustrate the lower bound $\Omega(\min\{k^2, nr\})$ predicted by Theorem 3.3.

The penalty function $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ depends on the norm of \mathbf{X}_0 . In practice the norm of the solution is not known beforehand, a weighted sum of norms is used instead. In Figure 6 we examine the performance of the weighted sum of norms penalty in recovery of a rank-1 PSD matrix, for different weights. We pick $\lambda = 0.20$ and $\lambda = 0.35$ for a randomly generated matrix \mathbf{X}_0 , and it can be seen that we get a reasonable result which is comparable to the performance of $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$.

In addition, we consider the *amount of error* in the recovery when the program fails. Figure 7 shows two curves below which we get a 90% percent failure, where for the green curve the normalized error threshold for




Figure 6: Performance of the recovery program minimizing $\text{tr}(\mathbf{X}) + \lambda \|\mathbf{X}\|_1$ with a PSD constraint, for $\lambda = 0.2$ (left) and $\lambda = 0.35$ (right).

declaring failure is 10^{-4} , and for the red curve it is a larger value of 0.05. We minimize $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ as the objective. We observe that when the recovery program has an error, it is very likely that this error is large, as the curves for 10^{-4} and 0.05 almost overlap. Thus, when the program fails, it fails badly. This observation agrees with intuition from similar problems in compressed sensing where sharp phase transition is observed.




Figure 7: 90% frequency of failure where the threshold of recovery is 10^{-4} for the green and 0.05 for the red curve. $\max\{\frac{\text{tr}(\mathbf{X})}{\text{tr}(\mathbf{X}_0)}, \frac{\|\mathbf{X}\|_1}{\|\mathbf{X}_0\|_1}\}$ is minimized subject to the PSD constraint and the measurements.

As a final comment, observe that, in Figures 5, 6 and 7 the required amount of measurements slowly increases even when n is large and $k^2 = 64$ is the dominant constant term. While this is consistent with our lower bound of $\Omega(k^2, n)$, the slow increase for constant k , can be explained by the fact that, as n gets larger, sparsity becomes the dominant structure and ℓ_1 minimization by itself requires $O(k^2 \log \frac{n}{k})$ measurements rather than $O(k^2)$. Hence for large n , the number of measurements can be expected to grow logarithmically in n .

7 Discussion

We have considered the problem of recovery of a simultaneously structured object from limited measurements. It is common in practice to combine known norm penalties corresponding to the individual structures (also

known as regularizers in statistics and machine learning applications), and minimize this combined objective in order to recover the object of interest. The common use of this approach motivated us to analyze its performance, in terms of the smallest number of generic measurements needed for correct recovery. We showed that, under a certain assumption on the norms involved, the combined penalty requires more generic measurements than one would expect based on the degrees of freedom of the desired object. Our lower bounds on the required number of measurements implies that the combined norm penalty cannot perform significantly better than the best individual norm.

These results raise several interesting questions, and lead to directions for future work. We briefly outline some of these directions, as well as connections to some related problems.

Other types of measurements, and sparse phase retrieval. Our current analysis is limited to random Gaussian measurements, and in some applications other kinds of measurements need to be considered. In the sparse phase retrieval problem, the measurements are given by $\langle \mathbf{a}_i \mathbf{a}_i^T, \mathbf{X} \rangle = b_i$, where \mathbf{a}_i is a random Gaussian vector and we wish to recover the underlying matrix $\mathbf{X}_0 = \mathbf{x}_0 \mathbf{x}_0^T$ where \mathbf{x}_0 is a k sparse vector. In this case, a different probabilistic argument might be necessary due to the product of random Gaussian variables that appear. In [31] the authors derive recovery guarantees for these measurements, but they do not consider signal sparsity, and their analysis does not immediately extend to the sparse case. The very recent paper [25] considers the sparse phase retrieval problem, and gives two results, assuming \mathbf{x}_0 has unit ℓ_2 norm and $m \ll n$: first, if $m \geq \mathcal{O}(\|\mathbf{x}_0\|_1^2 k \log n)$ then minimizing $\|\mathbf{X}\|_1 + \lambda \text{tr}(\mathbf{X})$ for suitable value of λ over the set of PSD matrices will exactly recover \mathbf{X}_0 with high probability. Second, they give a necessary condition on the number of measurements, $m'_0 = \min\left(\left(\frac{k}{4} - 1\right)^2, \frac{\max(\|\mathbf{x}_0\|_1^2 - k/2, 0)^2}{500 \log^2 n}\right)$, under which the recovery program fails to recover \mathbf{X}_0 with high probability whenever $m \leq m'_0$.

We should emphasize that the lower bound provided in [25] is directly comparable to our results. First, observe that since \mathbf{X}_0 is rank 1, and we are using ℓ_1 and nuclear norm subject to a PSD constraint, so the problem falls into Model (c) of Definition 3.2. Our result regarding this model is given in Theorem 3.3 which suggests that one needs at least $m'_0 = c \frac{\|\mathbf{x}_0\|_1^2}{k} \min\{k^2, n\}$ measurements. Assuming $m \ll n$, our lower bound takes a simpler form which is $m'_0 = c \|\mathbf{x}_0\|_1^2 k$.

Now, comparing our lower bound to the bound given by [25], we have

$$m''_0 \leq \frac{\max(\|\mathbf{x}_0\|_1^2 - k/2, 0)^2}{500 \log^2 n} \leq \frac{\|\mathbf{x}_0\|_1^4}{500 \log^2 n} \leq \frac{\|\mathbf{x}_0\|_1^2 k}{500 \log^2 n} \leq 500c \frac{m'_0}{\log^2(n)}.$$

where we use the fact that $\|\mathbf{x}_0\|_1 \leq \sqrt{k}$ for unit length and k sparse vector \mathbf{x}_0 . These simple operations suggest our lower bound is larger by a factor of $\frac{k \log^2(n)}{\|\mathbf{x}_0\|_1^2} \geq \log^2 n$ where we omit the constant term $500c$ for the sake of clarity. Overall, while the present paper and [25] analyze the same problem with different measurement operators, the results are consistent as logarithmic terms have relatively minor importance.

It is also of interest to study recovery properties of simultaneously structured models using other classes of measurements, for example cases where the measurement vectors $\{\mathbf{a}_i\}_{i=1}^m$ are binary or are sampled from the rows of a Discrete Fourier Transform matrix.

Quantifying recovery failure via error bounds. We observe from the recovery error plots shown in Figure 7 that whenever our recovery program fails, it fails with a significant recovery error. The figure shows two curves under which recovery fails with high probability, where failure is defined by the normalized error $\|\mathbf{X} - \mathbf{X}_0\|_F / \|\mathbf{X}_0\|_F$ being above 10^{-4} and 0.05. The two curves almost coincide. This observation leads to the question of whether we can characterize how large the error is with a high probability over the random measurements. A lower bound on the recovery error as a function of the number of problem parameters will be very insightful.

Defining new atoms for simultaneously structured models. Our results show that combinations of individual norms do not exhibit a strong recovery performance. On the other hand, the seminal paper [8]

proposes a remarkably general construction for an appropriate penalty given a set of atoms. Can we revisit a simultaneously structured recovery problem, and define new atoms that capture all structures at the same time? And can we obtain a new norm penalty induced by the convex hull of the atoms? Abstractly, the answer is yes, but such convex hulls may be hard to characterize, and the corresponding penalty may not be efficiently computable. It is interesting to find special cases where this construction can be carried out and results in a tractable problem.

Algorithms for minimizing combination of norms. Despite the limitation in their theoretical performance, in practice one may still need to solve convex relaxations that combine the different norms, i.e., problem (2.7). Consider the special case of sparse and low-rank matrix recovery. All corresponding optimization problems mentioned in Theorem 3.3 can be expressed as a semidefinite program and solved by standard solvers; for example, for the numerical experiments in Section 6 we used the interior-point solver SeDuMi [43] via the modeling environment CVX [42]. However, interior point methods do not scale for problems with tens of thousands of matrix entries, which are common in machine learning applications. One future research direction is to explore first-order methods, which have been successful in solving problems with a single structure (for example ℓ_1 or nuclear norm regularization alone). In particular, Alternating Directions Methods of Multipliers (ADMM) appears to be a promising candidate.

Characterizing the tightness of the lower bounds. The results provided in this paper are negative in nature, as we characterize the lower bounds on the required amount of measurements for mixed convex recovery problems. However, it would be interesting to see how much we can gain by making use of multiple norms and how tight are these lower bounds. In [44], authors investigate a specific simultaneous model where signal $\mathbf{x} \in \mathbb{R}^n$ is sparse in both time and frequency domains, i.e., \mathbf{x} and $\mathbf{D}\mathbf{x}$ are k_1, k_2 sparse respectively where \mathbf{D} is the Discrete Fourier Transform matrix. For recovery, the authors consider minimizing $\|\mathbf{x}\|_1 + \lambda\|\mathbf{D}\mathbf{x}\|_1$ subject to measurements. Intuitively, results of this paper would suggest the necessity of $\Omega(\min\{k_1, k_2\})$ measurements for successful recovery. On the other hand, best of the individual functions (ℓ_1 norms) will require $\Omega(\min\{k_1 \log \frac{n}{k_1}, k_2 \log \frac{n}{k_2}\})$ measurements. In [44], it is shown that the mixed approach will require as little as $\max\{k_1, k_2\} \log \log n$ under mild assumptions.

This shows that the mixed approach can result in a logarithmic improvement over the individual functions when $k_1 \approx k_2$ and the lower bound given by this paper can be achieved up to a $\log \log n$ factor.

Connection to Sparse PCA. The sparse PCA problem (see, e.g. [45, 46, 47]) seeks sparse principal components given a (possibly noisy) data matrix. Several formulations for this problem exist, and many algorithms have been proposed. In particular, a popular algorithm is the SDP relaxation proposed in [47], which is based on the following formulation.

For the first principal component to be sparse, we seek an $\mathbf{x} \in \mathbb{R}^n$ that maximizes $\mathbf{x}^T \mathbf{A} \mathbf{x}$ for a given data matrix \mathbf{A} , and minimizes $\|\mathbf{x}\|_0$. Similar to the sparse phase retrieval problem, this problem can be reformulated in terms of a rank-1, PSD matrix $\mathbf{X} = \mathbf{x}\mathbf{x}^T$ which is also row- and column-sparse. Thus we seek a simultaneously low-rank and sparse \mathbf{X} . This problem is different from the recovery problem studied in this paper, since we do not have m random measurements of \mathbf{X} . Yet, it will be interesting to connect this paper's results to the sparse PCA problem to potentially provide new insights for sparse PCA.

Acknowledgements. This work was supported in part by the National Science Foundation under grants CCF-0729203, CNS-0932428 and CCF-1018927, by the Office of Naval Research under the MURI grant N00014-08-1-0747, by Caltech's Lee Center for Advanced Networking, and by the National Science Foundation CAREER award ECCS-0847077. The work of Y. Eldar is supported in part by the Israel Science Foundation under Grant no. 170/10, in part by the Ollendorf Foundation, and in part by a Magnet grant Metro450 from the Israel Ministry of Industry and Trade.

References

- [1] E.J. Candès and T. Tao, “Decoding by linear programming,” *IEEE Trans. Inform. Theory*, 51 4203-4215.
- [2] D.L. Donoho, “Compressed sensing,” *IEEE Trans. Inform. Theory*, 52(4):1289-1306, 2006.
- [3] E.J. Candès, JK Romberg and T Tao, “Stable signal recovery from incomplete and inaccurate measurements”. *Comm. on Pure and Applied Math.* Vol. 59, Issue 8, pg 12071223, August 2006.
- [4] B. Recht, M. Fazel, P. Parrilo, “Guaranteed Minimum-Rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization”. *SIAM Review*, Vol 52, no 3, pages 471-501, 2010.
- [5] E.J. Candès and B. Recht, “Exact matrix completion via convex optimization,” *Found. of Comput. Math.*, 9 717-772.
- [6] V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky, “Latent Variable Graphical Model Selection via Convex Optimization”, *Annals of Statistics*.
- [7] E.J. Candès, X. Li, Y. Ma, J. Wright, “Robust Principal Component Analysis?”. *Journal of ACM* 58(1), 1-37.
- [8] V. Chandrasekaran, B. Recht, P. A. Parrilo, A. S. Willsky, “The Convex Geometry of Linear Inverse Problems”. *arXiv:1012.0621v3*.
- [9] A. Agarwal, S. Negahban, M. J. Wainwright, “Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions,” *Annals of Statistics*, Volume 40, Number 2 (2012), 1171-1197.
- [10] E.J. Candès, Y. Plan. “Tight oracle bounds for low-rank matrix recovery from a minimal number of random measurements,” *IEEE Transactions on Information Theory* 57(4), 2342-2359.
- [11] E.J. Candès, J. Romberg, T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information”. *IEEE Trans. Inform. Theory*, 52 489–509.
- [12] J. Wright, A. Ganesh, K. Min, Y. Ma, “Compressive Principal Component Pursuit”. *arXiv:1202.4596v1*.
- [13] M. Stojnic, F. Parvaresh, and B. Hassibi, On the reconstruction of block-sparse signals with an optimal number of measurements, *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 30753085, May 2010.
- [14] M. Yuan and Y. Lin, Model selection and estimation in regression with grouped variables, *J. Roy. Stat. Soc. Ser. B Stat. Methodol.*, vol. 68, no. 1, pp. 4967, 2006.
- [15] P. Sprechmann, I. Ramirez, G. Sapiro, Y.C. Eldar, “C-HiLasso: A Collaborative Hierarchical Sparse Modeling Framework”, *IEEE Transactions on Signal Processing*, vol.59, issue 9, pp.4183-4198, Sept. 2011.
- [16] M. Golbabaee and P. Vandergheynst, “Hyperspectral image compressed sensing via low-rank and joint-sparse matrix recovery”. *ICASSP* 2012.
- [17] Y. Shechtman and Y.C. Eldar and A. Szameit and M. Segev, “Sparsity-based sub-wavelength imaging with partially spatially incoherent light via quadratic compressed sensing”, *Optics Express*, 19:14807–14822 , 2011.
- [18] A. Beck, Y.C. Eldar, “Sparsity constrained nonlinear optimization: Optimality conditions and algorithms”, *arXiv:1203.4580v1*,
- [19] A. Szameit et. al., “Sparsity-based single-shot sub-wavelength coherent diffractive imaging”, *Nature Materials*.

- [20] A. Walther. “The question of phase retrieval in optics”. *Opt. Acta*, 10:41–49, 1963.
- [21] R.P. Millane, “Phase retrieval in crystallography and optics”. *J. Opt. Soc. Am. A* 7, 394-411 (1990).
- [22] R.W. Harrison, “Phase problem in crystallography”. *J. Opt. Soc. Am. A*, 10(5):1045–1055, 1993.
- [23] N. Hurt. “Phase retrieval and zero crossings”, Kluwer Academic Publishers, Norwell, MA, 1989.
- [24] H. Ohlsson, A. Y. Yang, R. Dong, S. S. Sastry, “Compressive phase retrieval from squared output measurements via semidefinite programming”, arXiv:1111.6323v3, March 2012.
- [25] X. Li, V. Voroninski, “Sparse Signal Recovery from Quadratic Measurements via Convex Programming,” arXiv:1209.4785.
- [26] K. Jaganathan, S. Oymak, and B. Hassibi, “Recovery of sparse 1-D signals from the magnitudes of their Fourier transform”, arXiv:1206.1405v1, June 2012.
- [27] Y.M. Lu and M. Vetterli, “Sparse Spectral Factorization: Unicity and Reconstruction Algorithms”. *Acoustics, Speech and Signal Processing (ICASSP)*, 2011 IEEE International Conference on, pp. 5976-5979, 22-27 May 2011.
- [28] Y. Shechtman and A. Beck and Y.C. Eldar, “Efficient Phase Retrieval of Sparse Signals”. *IEEEI* 2012.
- [29] Y.C. Eldar and S. Mendelson ”Phase Retrieval: Stability and Recovery Guarantees”, Arxiv:1211.0872, Nov. 2012.
- [30] E.J. Candès, Y.C. Eldar, T. Strohmer and V. Voroninski, “Phase retrieval via matrix completion”, arXiv:1109.0573, Sep. 2011.
- [31] E.J. Candès, T. Strohmer, V. Voroninski. “PhaseLift: exact and stable signal recovery from magnitude measurements via convex programming”. to appear in *Communications on Pure and Applied Mathematics*.
- [32] E.J. Candès, X. Li, “Solving Quadratic Equations via PhaseLift when There Are About As Many Equations As Unknowns”, arXiv:1208.6247, Aug. 2012.
- [33] B. Alexeev, A.S. Bandeira, M. Fickus, D.G. Mixon, “Phase retrieval with polarization”, arXiv:1210.7752.
- [34] J.R. Fienup, “Phase retrieval algorithms: a comparison”, *Applied Optics* 21, 2758-2769, 1982.
- [35] R.W. Gerchberg and W.O. Saxton, “Phase retrieval by iterated projections”, *Optik* 35, 237, 1972.
- [36] E.J. Candès and B. Recht, “Simple Bounds for Recovering Low-complexity Models”. arXiv:1106.1474v2.
- [37] G.A. Watson, “Characterization of the subdifferential of some matrix norms”. *Linear Algebra and its Appl.* 170, 33–45 (1992).
- [38] S. Boyd and L. Vandenberghe, “Convex Optimization”. Cambridge University Press, 2004.
- [39] D. Bertsekas with A. Nedic and A.E. Ozdaglar, “Convex Analysis and Optimization” Athena Scientific, 2003.
- [40] R.T. Rockafellar, R. J-B Wets, “Variational Analysis”. Springer, 2004.
- [41] R.T. Rockafellar, “Convex Analysis”. Princeton University Press, 1997.
- [42] CVX Research, Inc. “CVX: Matlab software for disciplined convex programming”, version 2.0 beta. <http://cvxr.com/cvx>, September 2012.
- [43] J.F. Sturm, “Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones,” 1998.

- [44] S. Oymak and B. Hassibi, “Recovering Jointly Sparse Signals via Joint Basis Pursuit”. arXiv:1202.3531v1.
- [45] M. Journée, Y. Nesterov, P. Richtárik, and R Sepulchre, “Generalized Power Method for Sparse Principal Component Analysis”. *Journal of Machine Learning Research* 11 (2010) 517–553.
- [46] H. Zou, T. Hastie, and R. Tibshirani, “Sparse Principal Component Analysis”, *Journal of computational and graphical statistics* 15 (2), 265-286.
- [47] A. d’Aspremont, F. Bach, L. El Ghaoui. “Optimal Solutions for Sparse Principal Component Analysis”, *Journal of Machine Learning Research*, 9(Jul):12691294, 2008.
- [48] J.J. Moreau, “Décomposition orthogonale d’un espace Hilbertien selon deux cones mutuellement polaires”. *C. R. Acad. Sci.*, volume 255, pages 238–240, 1962.
- [49] Y. Gordon, “On Milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n ”. *Geometric aspects of functional analysis*, *Isr. Semin. 1986-87*, *Lect. Notes Math.* 1317, 84-106. (1988).
- [50] R. Vershynin, “Introduction to the non-asymptotic analysis of random matrices”. arXiv:1011.3027v7.

APPENDIX

A Properties of Cones

In this appendix, we state some results regarding cones which are used in the proof of general recovery. Recall the definitions of polar and dual cones from Section 2.

Theorem A.1 (Moreau's decomposition theorem, [48]) *Let \mathcal{C} be a closed and convex cone in \mathbb{R}^n . Then, for any $\mathbf{x} \in \mathbb{R}^n$, we have*

- $\mathbf{x} = \mathcal{P}_{\mathcal{C}}(\mathbf{x}) + \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x})$.
- $\langle \mathcal{P}_{\mathcal{C}}(\mathbf{x}), \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x}) \rangle = 0$.

Lemma A.1 (Projection is nonexpansive) *Let $\mathcal{C} \in \mathbb{R}^n$ be a closed and convex set and $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ be vectors. Then,*

$$\|\mathcal{P}_{\mathcal{C}}(\mathbf{a}) - \mathcal{P}_{\mathcal{C}}(\mathbf{b})\|_2 \leq \|\mathbf{a} - \mathbf{b}\|_2. \quad (\text{A.1})$$

Corollary A.1 *Let \mathcal{C} be a closed convex cone and \mathbf{a}, \mathbf{b} be vectors satisfying $\mathcal{P}_{\mathcal{C}}(\mathbf{a} - \mathbf{b}) = 0$. Then*

$$\|\mathbf{b}\|_2 \geq \|\mathcal{P}_{\mathcal{C}}(\mathbf{a})\|_2. \quad (\text{A.2})$$

Proof. Using Lemma A.1, we have $\|\mathcal{P}_{\mathcal{C}}(\mathbf{a})\|_2 = \|\mathcal{P}_{\mathcal{C}}(\mathbf{a}) - \mathcal{P}_{\mathcal{C}}(\mathbf{a} - \mathbf{b})\|_2 \leq \|\mathbf{b}\|_2$. ■

The unit sphere in \mathbb{R}^n will be denoted by \mathcal{S}^{n-1} for the following theorems.

Theorem A.2 (Escape through a mesh, [49]) *For a given set $\mathcal{D} \in \mathcal{S}^{n-1}$, define the Gaussian width as*

$$\omega(\mathcal{D}) = \mathbb{E} \left[\sup_{\mathbf{x} \in \mathcal{D}} \langle \mathbf{x}, \mathbf{g} \rangle \right],$$

in which $\mathbf{g} \in \mathbb{R}^n$ has i.i.d. standard Gaussian entries. Given m , let $d = \sqrt{n-m} - \frac{1}{4\sqrt{n-m}}$. Provided that $\omega(\mathcal{D}) \leq d$ a random m -dimensional subspace which is uniformly drawn w.r.t. Haar measure will have no intersection with \mathcal{D} with probability at least

$$1 - 3.5 \exp(-(d - \omega(\mathcal{D}))^2). \quad (\text{A.3})$$

Theorem A.3 *Consider a random Gaussian map $\mathcal{G} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ with i.i.d. entries and the corresponding adjoint operator \mathcal{G}^* . Let \mathcal{C} be a closed and convex cone and recalling Definition 3.1, let*

$$\zeta(\mathcal{C}) := 1 - \eta(\mathcal{C}^\circ), \quad \gamma(\mathcal{C}) := \frac{2\eta(\mathcal{C})}{1 - \eta(\mathcal{C}^\circ)}.$$

Then, if $m \leq \frac{7\zeta(\mathcal{C})}{16}n$, with probability at least $1 - 6 \exp(-(\frac{\zeta(\mathcal{C})}{4})^2 n)$, for all $\mathbf{z} \in \mathbb{R}^n$ we have

$$\|\mathcal{G}^*(\mathbf{z})\|_2 \leq \gamma(\mathcal{C}) \|\mathcal{P}_{\mathcal{C}}(\mathcal{G}^*(\mathbf{z}))\|_2. \quad (\text{A.4})$$

Proof. For notational simplicity, let $\zeta = \zeta(\mathcal{C})$ and $\gamma = \gamma(\mathcal{C})$. Consider the set

$$\mathcal{D} = \{\mathbf{x} \in \mathcal{S}^{n-1} : \|\mathbf{x}\|_2 \geq \gamma \|\mathcal{P}_{\mathcal{C}}(\mathbf{x})\|_2\}.$$

and we are going to show that with high probability, the range of \mathcal{G}^* misses \mathcal{D} . Using Theorem A.1, for any $\mathbf{x} \in \mathcal{D}$, we may write

$$\begin{aligned} \langle \mathbf{x}, \mathbf{g} \rangle &= \langle \mathcal{P}_{\mathcal{C}}(\mathbf{x}) + \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x}), \mathcal{P}_{\mathcal{C}}(\mathbf{g}) + \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g}) \rangle \\ &\leq \langle \mathcal{P}_{\mathcal{C}}(\mathbf{x}), \mathcal{P}_{\mathcal{C}}(\mathbf{g}) \rangle + \langle \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x}), \mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g}) \rangle \\ &\leq \|\mathcal{P}_{\mathcal{C}}(\mathbf{x})\|_2 \|\mathcal{P}_{\mathcal{C}}(\mathbf{g})\|_2 + \|\mathcal{P}_{\mathcal{C}^\circ}(\mathbf{x})\|_2 \|\mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g})\|_2 \\ &\leq \gamma^{-1} \|\mathcal{P}_{\mathcal{C}}(\mathbf{g})\|_2 + \|\mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g})\|_2 \end{aligned} \quad (\text{A.5})$$

where in (A.5) we used the fact that elements of \mathcal{C} and \mathcal{C}° have nonpositive inner products and $\|\mathcal{P}_{\mathcal{C}}(\mathbf{x})\|_2 \leq \|\mathbf{x}\|_2$ is by Lemma A.1. Hence, from the definition of Gaussian width,

$$\begin{aligned}\omega(\mathcal{D}) &= \mathbb{E} \left[\sup_{\mathbf{x} \in \mathcal{D}} \langle \mathbf{x}, \mathbf{g} \rangle \right] \leq \gamma^{-1} \mathbb{E} [\|\mathcal{P}_{\mathcal{C}}(\mathbf{g})\|_2] + \mathbb{E} [\|\mathcal{P}_{\mathcal{C}^\circ}(\mathbf{g})\|_2] \\ &= \sqrt{n}(\gamma^{-1}\eta(\mathcal{C}) + \eta(\mathcal{C}^\circ)) = \frac{2-\zeta}{2}\sqrt{n}.\end{aligned}$$

Now, whenever

$$m \leq \frac{7\zeta}{16}n \leq (1 - (\frac{4-\zeta}{4})^2)n = m', \quad (\text{A.6})$$

we have

$$(\sqrt{n-m} - \omega(\mathcal{D}) - \frac{1}{4\sqrt{n-m}})^2 \geq (\sqrt{n-m} - \omega(\mathcal{D}))^2 - \frac{1}{2} \geq (\frac{\zeta}{4})^2 n - \frac{1}{2}. \quad (\text{A.7})$$

Now, using Theorem A.2, the range space of \mathcal{G}^* will miss the undesired set \mathcal{D} with probability at least $1 - 3.5 \exp(-(\frac{\zeta}{4})^2 n + \frac{1}{2}) \geq 1 - 6 \exp(-(\frac{\zeta}{4})^2 n)$. \blacksquare

Corollary A.2 Consider the cones \mathbb{S}^n and \mathbb{S}_+^n in the space $\mathbb{R}^{n \times n}$. For all positive integers n , we have

- $\zeta(\mathbb{S}^n) \geq \frac{1}{4}$ and $\gamma(\mathbb{S}^n) \leq 7$.
- $\zeta(\mathbb{S}_+^n) \geq \frac{1}{8}$ and $\gamma(\mathbb{S}_+^n) \leq 11$.

Proof. Let \mathbf{G} be an $n \times n$ matrix with i.i.d. standard normal entries. Set of symmetric matrices \mathbb{S}^n is an $\frac{n(n+1)}{2}$ dimensional subspace of $\mathbb{R}^{n \times n}$. Hence, $\mathbb{E} \|\mathcal{P}_{\mathbb{S}^n}(\mathbf{G})\|_F^2 = \frac{n(n+1)}{2}$ and $\mathbb{E} \|\mathcal{P}_{(\mathbb{S}^n)^\circ}(\mathbf{G})\|_F^2 = \frac{n(n-1)}{2}$.

$$\zeta(\mathbb{S}^n) = 1 - \frac{\mathbb{E} [\|\mathcal{P}_{(\mathbb{S}^n)^\circ}(\mathbf{G})\|_F]}{n} \geq 1 - \frac{\sqrt{\mathbb{E} [\|\mathcal{P}_{(\mathbb{S}^n)^\circ}(\mathbf{G})\|_F^2]}}{n} > 1 - \sqrt{\frac{n-1}{2n}} > 1 - \frac{1}{\sqrt{2}}. \quad (\text{A.8})$$

Similarly, $\gamma(\mathbb{S}^n) = \frac{2\eta(\mathbb{S}^n)}{\zeta(\mathbb{S}^n)} < 7\eta(\mathbb{S}^n) \leq 7$.

To prove the second statement, observe that projection of a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ onto \mathbb{S}_+^n is obtained by first projecting \mathbf{A} onto \mathbb{S}^n and then taking the matrix induced by the positive eigenvalues of $\mathcal{P}_{\mathbb{S}^n}(\mathbf{A})$. Since, \mathbf{G} and $-\mathbf{G}$ are identically distributed and \mathbb{S}_+^n is a self dual cone, $\mathcal{P}_{\mathbb{S}_+^n}(\mathbf{G})$ is identically distributed as $-\mathcal{P}_{\mathbb{S}_-^n}(\mathbf{G})$ where $\mathbb{S}_-^n = (\mathbb{S}_+^n)^\circ$ stands for negative semidefinite matrices. Hence,

$$\mathbb{E} \|\mathcal{P}_{\mathbb{S}_+^n}(\mathbf{G})\|_F^2 = \frac{\mathbb{E} \|\mathcal{P}_{\mathbb{S}^n}(\mathbf{G})\|_F^2}{2} = \frac{n(n+1)}{4}, \quad \mathbb{E} \|\mathcal{P}_{(\mathbb{S}_+^n)^\circ}(\mathbf{G})\|_F^2 = \frac{n(3n-1)}{4}. \quad (\text{A.9})$$

Consequently,

$$\zeta(\mathbb{S}_+^n) \geq 1 - \frac{\sqrt{\mathbb{E} [\|\mathcal{P}_{(\mathbb{S}_+^n)^\circ}(\mathbf{G})\|_F^2]}}{n} > 1 - \sqrt{\frac{3n-1}{4n}} > 1 - \frac{\sqrt{3}}{2}. \quad (\text{A.10})$$

$$\gamma(\mathbb{S}_+^n) = \frac{2\eta(\mathbb{S}_+^n)}{\zeta(\mathbb{S}_+^n)} < 15\eta(\mathbb{S}_+^n) \leq 15\sqrt{\frac{n(n+1)}{4n^2}} \leq \frac{15}{\sqrt{2}} < 11. \quad \blacksquare$$

B Properties of Norms

Lemma B.1 Given a norm $\|\cdot\|$, denote by L the Lipschitz constant of this norm. For any \mathbf{x} , we have

$$L = \sup_{\|\mathbf{y}\|_2 \leq 1} \|\mathbf{y}\| = \sup_{\|\mathbf{z}\|^* \leq 1} \|\mathbf{z}\|_2 \geq \sup_{\mathbf{g} \in \partial\|\mathbf{x}\|} \|\mathbf{g}\|_2. \quad (\text{B.1})$$

Proof. Let $M = \sup_{\|\mathbf{y}\|_2 \leq 1} \|\mathbf{y}\|$ and $C = \sup_{\|\mathbf{z}\|^* \leq 1} \|\mathbf{z}\|_2$. We have the followings.

- $M = L$ because on one hand for all \mathbf{x} we have $\|\mathbf{x}\| = \|\mathbf{x}\| - \|0\| \leq L\|\mathbf{x}\|_2$ which implies $M \leq L$, and on the other hand $\|\|\mathbf{x}\| - \|\mathbf{y}\|\| \leq \|\mathbf{x} - \mathbf{y}\| \leq M\|\mathbf{x} - \mathbf{y}\|_2$ which holds for all \mathbf{x} and \mathbf{y} and implies $L \leq M$.
- $M \leq C$ because for any \mathbf{x} we have

$$\|\mathbf{x}\| = \sup_{\|\mathbf{z}\|^* \leq 1} \langle \mathbf{x}, \mathbf{z} \rangle \leq \sup_{\|\mathbf{z}\|^* \leq 1} \|\mathbf{z}\|_2 \|\mathbf{x}\|_2.$$

- In addition, let $\mathbf{z}_0 = \arg \sup_{\|\mathbf{z}\|^* \leq 1} \|\mathbf{z}\|_2$. Then, we have

$$C^2 = \|\mathbf{z}_0\|_2^2 = \langle \mathbf{z}_0, \mathbf{z}_0 \rangle = \sup_{\|\mathbf{z}\|^* \leq 1} \langle \mathbf{z}, \mathbf{z}_0 \rangle = \|\mathbf{z}_0\|,$$

which gives $C = \frac{\|\mathbf{z}_0\|}{\|\mathbf{z}_0\|_2} \leq L$.

Overall, we obtain $L = M = C$. Finally, since $\partial\|\mathbf{x}\| \subseteq \{\mathbf{z} : \|\mathbf{z}\|^* \leq 1\}$, we get the last inequality in (B.1). ■

B.1 Norms in Sparse and Low-rank Model

B.1.1 Relevant notation for the proofs

Let $[k]$ denote the set $\{1, 2, \dots, k\}$. Let S_c, S_r denote the indexes of the nonzero columns and rows of \mathbf{X}_0 so that nonzero entries of \mathbf{X}_0 lies on $S_r \times S_c$ submatrix. $\mathcal{S}_c, \mathcal{S}_r$ denotes the k_1, k_2 dimensional subspaces of vectors whose nonzero entries lie on S_c and S_r respectively.

Let \mathbf{X}_0 have singular value decomposition $\mathbf{U}\Sigma\mathbf{V}^T$ such that $\Sigma \in \mathbb{R}^{r \times r}$ and columns of \mathbf{U}, \mathbf{V} lies on $\mathcal{S}_c, \mathcal{S}_r$ respectively.

B.1.2 Proof of Lemma 5.1

Proof. Observe that $T_c = \mathbb{R}^n \times \mathcal{S}_c$ and $T_r = \mathcal{S}_r \times \mathbb{R}^n$ hence $T_c \cap T_r$ is the set of matrices that lies on $S_r \times S_c$. Hence, $\mathbf{E}_\star = \mathbf{U}\mathbf{V}^T \in T_c \cap T_r$. Similarly, \mathbf{E}_c and \mathbf{E}_r are the matrices obtained by scaling columns and rows of \mathbf{X}_0 to have unit size. As a result, they lie on $S_r \times S_c$ and $T_c \cap T_r$. $\mathbf{E}_\star \in T_\star$ by definition.

Next, we may write $\mathbf{E}_c = \mathbf{X}_0 \mathbf{D}_c$ where \mathbf{D}_c is the scaling nonnegative diagonal matrix. Consequently, \mathbf{E}_c lies on the range space of \mathbf{X}_0 and belongs to T_\star . This follows from definition of T_\star in Section 2 and the fact that $(\mathbf{I} - \mathbf{U}\mathbf{U}^T)\mathbf{E}_c = 0$.

In the exact same way, $\mathbf{E}_r = \mathbf{D}_r \mathbf{X}_0$ for some nonnegative diagonal \mathbf{D}_r and lies on the range space of \mathbf{X}^T and hence lies on T_\star . Consequently, $\mathbf{E}_\star, \mathbf{E}_c, \mathbf{E}_r$ lies on $T_c \cap T_r \cap T_\star$.

Now, consider

$$\langle \mathbf{E}_c, \mathbf{E}_\star \rangle = \langle \mathbf{X}_0 \mathbf{D}_c, \mathbf{U}\mathbf{V}^T \rangle = \text{tr}(\mathbf{V}\mathbf{U}^T \mathbf{U}\Sigma\mathbf{V}^T \mathbf{D}_c) = \text{tr}(\mathbf{V}\Sigma\mathbf{V}^T \mathbf{D}_c) \geq 0.$$

since both $\mathbf{V}\Sigma\mathbf{V}^T$ and \mathbf{D}_c are positive semidefinite matrices. In the exact same way, we have $\langle \mathbf{E}_c, \mathbf{E}_\star \rangle \geq 0$. Finally,

$$\langle \mathbf{E}_c, \mathbf{E}_r \rangle = \langle \mathbf{X}_0 \mathbf{D}_c, \mathbf{D}_r \mathbf{X}_0 \rangle = \text{tr}(\mathbf{D}_c \mathbf{X}_0^T \mathbf{D}_r \mathbf{X}_0) \geq 0, \quad (\text{B.2})$$

since both \mathbf{D}_c and $\mathbf{X}_0^T \mathbf{D}_r \mathbf{X}_0$ are PSD matrices. Overall, pairwise inner products of $\mathbf{E}_r, \mathbf{E}_c, \mathbf{E}_\star$ are nonnegative.

Now, we proceed with the remaining statements where \mathbf{X}_0 is rank one and we deal with ℓ_1 and nuclear norms. Since \mathbf{X}_0 is rank 1, we immediately have $\mathbf{E}_\star = \mathbf{u}\mathbf{v}^T \in T_1$ since nonzero locations of \mathbf{X}_0 and \mathbf{E}_\star are same. Observe that \mathbf{E}_1 is not necessarily inside T_\star but both $\mathbf{u}\mathbf{u}^T \mathbf{E}_1$ and $\mathbf{E}_1 \mathbf{v}\mathbf{v}^T$ is inside T_\star . Consequently, we have

$$\|\mathcal{P}_{T_\star}(\mathbf{E}_1)\|_F^2 = \|\mathbf{u}\mathbf{u}^T \mathbf{E}_1\|_F^2 + \|\mathbf{E}_1 \mathbf{v}\mathbf{v}^T\|_F^2 - \|\mathbf{u}\mathbf{u}^T \mathbf{E}_1 \mathbf{v}\mathbf{v}^T\|_F^2 \geq \max\{\|\mathbf{u}\mathbf{u}^T \mathbf{E}_1\|_F^2, \|\mathbf{E}_1 \mathbf{v}\mathbf{v}^T\|_F^2\}, \quad (\text{B.3})$$

where

$$\|\mathbf{u}\mathbf{u}^T \mathbf{E}_1\|_F = \|\mathbf{u}\mathbf{u}^T \text{sgn}(\mathbf{u}) \text{sgn}(\mathbf{v})^T\|_F = \|\mathbf{u}\|_1 \sqrt{k_2} = \frac{\|\mathbf{u}\|_1}{\sqrt{k_1}} \|\mathbf{E}_1\|_F, \quad (\text{B.4})$$

and similarly $\|\mathbf{E}_1 \mathbf{v}\mathbf{v}^T\|_F = \frac{\|\mathbf{v}\|_1}{\sqrt{k_2}} \|\mathbf{E}_1\|_F$. Hence,

$$\|\mathcal{P}_{T_\star}(\mathbf{E}_1)\|_F \geq \max\left\{\frac{\|\mathbf{u}\|_1}{\sqrt{k_1}}, \frac{\|\mathbf{v}\|_1}{\sqrt{k_2}}\right\} \|\mathbf{E}_1\|_F. \quad (\text{B.5})$$

Finally, since $\mathbf{E}_\star \in T_\star \cap T_1$, we have

$$\langle \mathcal{P}_{T_\star \cap T_1}(\mathbf{E}_1), \mathcal{P}_{T_\star \cap T_1}(\mathbf{E}_\star) \rangle = \langle \mathbf{E}_1, \mathcal{P}_{T_\star \cap T_1}(\mathbf{E}_\star) \rangle = \langle \mathbf{E}_1, \mathbf{E}_\star \rangle \quad (\text{B.6})$$

$$= \left\langle \text{sgn}(\mathbf{u}) \text{sgn}(\mathbf{v})^T, \mathbf{u}\mathbf{v}^T \right\rangle = \|\mathbf{u}\|_1 \|\mathbf{v}\|_1 \geq 0. \quad (\text{B.7})$$

■

Remark B.1 Observe that since \mathbf{u}, \mathbf{v} are unit length and k_1, k_2 sparse, we always have

$$\max\left\{\frac{\|\mathbf{u}\|_1}{\sqrt{k_1}}, \frac{\|\mathbf{v}\|_1}{\sqrt{k_2}}\right\} \leq 1,$$

where equality is achieved when $\sqrt{k_1}\mathbf{u}$ or $\sqrt{k_2}\mathbf{v}$ are sign vectors; i.e. nonzero entries are -1 or 1 . In general, we may let $\mathbf{X}_0 = \mathbf{a}\mathbf{b}^T$ where \mathbf{a}, \mathbf{b} are k_1, k_2 sparse and their nonzero entries are independent and identically distributed as random variables x_1, x_2 respectively. WLOG, let us consider \mathbf{a} . Assuming x_1 is zero-mean with finite fourth moment, for sufficiently large k_1 , with high probability, we would have,

$$\frac{\|\mathbf{u}\|_1}{\sqrt{k_1}} = \frac{\|\mathbf{a}\|_1}{\|\mathbf{a}\|_2} \sqrt{k_1} \approx \frac{\mathbb{E}|x_1|}{\sqrt{\mathbb{E}[x_1^2]}}. \quad (\text{B.8})$$

This follows from the fact that $\mathbb{E}[\|\mathbf{a}\|_1] = k_1 \mathbb{E}|x_1|$ and $\mathbb{E}[\|\mathbf{a}\|_2^2] = k_1 \mathbb{E}[x_1^2]$ and hence $\|\mathbf{a}\|_1$ and $\|\mathbf{a}\|_2$ will concentrate around $k_1 \mathbb{E}|x_1|$ and $\sqrt{k_1} \mathbb{E}[x_1^2]$ as k_1 grows. For example, when $x_1 \sim \mathcal{N}(0, 1)$, we have $\mathbb{E}[|x_1|] = \sqrt{\frac{2}{\pi}}$ and $\mathbb{E}[x_1^2] = 1$, hence $\frac{\|\mathbf{u}\|_1}{\sqrt{k_1}} \approx \sqrt{\frac{2}{\pi}}$. Overall, it is reasonable to consider $\frac{\|\mathbf{u}\|_1}{\sqrt{k_1}}$ as approximately constant as right hand side of (B.8) is only a function of the distribution of x_1 . The identical argument will apply for \mathbf{v} .

B.1.3 Results on positive semidefinite constraint

Lemma B.2 Assume $\mathbf{X}, \mathbf{Y} \in \mathbb{S}_+^n$ have eigenvalue decompositions $\mathbf{X} = \sum_{i=1}^{\text{rank}(\mathbf{X})} \sigma_i \mathbf{u}_i \mathbf{u}_i^T$ and $\mathbf{Y} = \sum_{i=1}^{\text{rank}(\mathbf{Y})} c_i \mathbf{v}_i \mathbf{v}_i^T$. Further, assume $\langle \mathbf{Y}, \mathbf{X} \rangle = 0$. Then, $\mathbf{U}^T \mathbf{Y} = 0$ where $\mathbf{U} = [\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_{\text{rank}(\mathbf{X})}]$.

Proof. Observe that,

$$\langle \mathbf{Y}, \mathbf{X} \rangle = \sum_{i=1}^{\text{rank}(\mathbf{X})} \sum_{j=1}^{\text{rank}(\mathbf{Y})} \sigma_i c_j |\mathbf{u}_i^T \mathbf{v}_j|^2. \quad (\text{B.9})$$

Since $\sigma_i, c_j > 0$, right hand side is 0 if and only if $\mathbf{u}_i^T \mathbf{v}_j = 0$ for all i, j . Hence, the result follows. ■

Lemma B.3 Assume $\mathbf{X}_0 \in \mathbb{S}_+^n$ so that in Section B.1.1, $S_c = S_r$, $T_c = T_r$, $k_1 = k_2 = k$ and $\mathbf{U} = \mathbf{V}$. Let $\mathbf{U}_1 \in \mathbb{R}^{n \times (k-r)}$ and $\mathbf{U}_2 \in \mathbb{R}^{n \times (n-k)}$ be such that $[\mathbf{U} \ \mathbf{U}_1]$ and $[\mathbf{U} \ \mathbf{U}_1 \ \mathbf{U}_2]$ be orthonormal bases over S_c and \mathbb{R}^n respectively. Also call $S_\star = T_\star \cap \mathbb{S}^n$ and let,

$$\mathcal{Y} = \{\mathbf{Y} | \mathbf{Y} \in (\mathbb{S}_+^n)^*, \ \langle \mathbf{Y}, \mathbf{X}_0 \rangle = 0\}. \quad (\text{B.10})$$

Then, the following statements hold.

- $S_\star = \text{span}(\mathcal{Y})^\perp$. Hence, recalling (2.6), $\mathcal{R} = T_\cap \cap S_\star$.
- Based on Definition 3.2, we have $\mathcal{R} = T_c \cap T_r \cap T_\star \cap \mathbb{S}^n$ for “PSD, arbitrary rank” model. Then, $\mathbf{E}_\star \in \mathcal{R}$, $\frac{\|\mathcal{P}_\mathcal{R}(\mathbf{E}_c)\|_F}{\|\mathbf{E}_c\|_F} = \frac{\|\mathcal{P}_\mathcal{R}(\mathbf{E}_r)\|_F}{\|\mathbf{E}_r\|_F} \geq \frac{1}{\sqrt{2}}$.
- Let \mathbf{X}_0 be rank 1, i.e. $\mathbf{X}_0 = \mathbf{x}_0 \mathbf{x}_0^T$. Then, $\mathcal{R} = T_1 \cap T_\star \cap \mathbb{S}^n$ for “PSD, rank 1” model. Then, $\mathbf{E}_\star \in \mathcal{R}$ and $\frac{\|\mathcal{P}_\mathcal{R}(\mathbf{E}_1)\|_F}{\|\mathbf{E}_1\|_F} \geq \frac{\|\mathbf{x}_0\|_1}{\sqrt{k}\|\mathbf{x}_0\|_2}$.

Proof. The dual of \mathbb{S}_+^n with respect to $\mathbb{R}^{n \times n}$ is the set sum of \mathbb{S}_+^n and Skew^n where Skew^n is the set of skew-symmetric matrices. Now, assume, $\mathbf{Y} \in \mathcal{Y}$ and $\mathbf{X} \in S_\star$. Then, $\langle \mathbf{Y}, \mathbf{X} \rangle = \langle \frac{\mathbf{Z}}{2}, \mathbf{X} \rangle$ where $\mathbf{Z} = \mathbf{Y} + \mathbf{Y}^T \in \mathbb{S}_+^n$ and $\langle \mathbf{Z}, \mathbf{X}_0 \rangle = 0$. Since \mathbf{X}_0, \mathbf{Z} are both PSD, applying Lemma B.2, we have $\mathbf{U}^T \mathbf{Z} = 0$ hence $(\mathbf{I} - \mathbf{U}\mathbf{U}^T)\mathbf{Z}(\mathbf{I} - \mathbf{U}\mathbf{U}^T) = \mathbf{Z}$ which means $\mathbf{Z} \in T_\star^\perp$. Hence, $\langle \mathbf{Z}, \mathbf{X} \rangle = \langle \mathbf{Y}, \mathbf{X} \rangle = 0$ as $\mathbf{X} \in S_\star \subset T_\star$. Hence, $\text{span}(\mathcal{Y}) \subseteq S_\star^\perp$.

On the other hand, $S_\star^\perp = T_\star^\perp + (\mathbb{S}^n)^\perp = T_\star^\perp + \text{Skew}^n$. Let $\mathbf{Y} \in S_\star^\perp$ and $\mathbf{Z} = \frac{\mathbf{Y} + \mathbf{Y}^T}{2} \in T_\star^\perp \cap \mathbb{S}^n$. Observe that $\mathbf{Y} - \mathbf{Z} \in \text{Skew}^n \in \mathcal{Y}$. Let \mathbf{Z} has eigenvalue decomposition $\mathbf{Z} = \sum_i \lambda_i \mathbf{z}_i \mathbf{z}_i^T$. Under proper unitary rotation, $T_\star^\perp \cap \mathbb{S}^n$ is equivalent to a set of matrices that are symmetric and supported over an $(n-r) \times (n-r)$ submatrix. Hence, eigenvectors also satisfy $\mathbf{z}_i \mathbf{z}_i^T \in T_\star^\perp \cap \mathbb{S}^n$. Now, observe that $\mathbf{z}_i \mathbf{z}_i^T \in (\mathbb{S}_+^n)^\star$ and $\langle \mathbf{z}_i \mathbf{z}_i^T, \mathbf{X}_0 \rangle = 0$. Hence, $\mathbf{z}_i \mathbf{z}_i^T \in \mathcal{Y}$ which implies $\mathbf{Z}, \mathbf{Y} \in \text{span}(\mathcal{Y})$. Overall, $S_\star^\perp \subseteq \text{span}(\mathcal{Y})$. Combined with the previous result we have $S_\star^\perp = \text{span}(\mathcal{Y})$.

For the second statement, $T_\cap = T_\star \cap T_c \cap T_r$ hence $\mathcal{R} = T_\star \cap T_c \cap T_r \cap \mathbb{S}^n$. Now, recalling Lemma 5.1, observe that we already know $\mathbf{E}_\star \in T_\cap$ where $T_\cap = T_c \cap T_r \cap T_\star$. Since \mathbf{E}_\star is also symmetric, $\mathbf{E}_\star \in \mathcal{R}$. Similarly, $\mathbf{E}_c \in T_\cap$, $\langle \mathbf{E}_c, \mathbf{E}_r \rangle \geq 0$ and $\|\mathcal{P}_\mathcal{R}(\mathbf{E}_c)\|_F = \frac{\|\mathbf{E}_c + \mathbf{E}_r\|_F}{2} \geq \frac{\|\mathbf{E}_c\|_F}{\sqrt{2}}$. Similar result is true for \mathbf{E}_r .

For third statement, $T_\cap = T_1 \cap T_\star$ and $\mathcal{R} = T_1 \cap T_\star \cap \mathbb{S}^n$. Now, observe that when \mathbf{X}_0 is rank 1, $T_1 = T_c \cap T_r$ hence $\mathbf{E}_\star \in \mathcal{R}$. Secondly, $\mathbf{E}_1 = \text{sgn}(\mathbf{x}_0) \text{sgn}(\mathbf{x}_0)^T \in \mathbb{S}$ and $\mathcal{P}_{T_\cap}(\mathbf{E}_1) \in \mathbb{S}$ as well since $\mathbf{X} \in T_\cap \implies \mathbf{X}^T \in T_\cap$.

Then, $\mathcal{P}_{T_\cap}(\mathbf{E}_1) \in \mathbb{S} \cap T_\cap = \mathcal{R}$. Finally, $\mathbf{E}_1 - \mathcal{P}_{T_\cap}(\mathbf{E}_1) \in T_\cap^\perp \subseteq \mathcal{R}^\perp$ as $\mathcal{R} \subseteq T_\cap$. This implies $\mathcal{P}_{T_\cap}(\mathbf{E}_1) = \mathcal{P}_\mathcal{R}(\mathbf{E}_1)$ when combined with $\mathcal{P}_{T_\cap}(\mathbf{E}_1) \in \mathcal{R}$. Hence using Lemma 5.1,

$$\|\mathcal{P}_\mathcal{R}(\mathbf{E}_1)\|_F = \|\mathcal{P}_{T_\cap}(\mathbf{E}_1)\|_F \geq \frac{\|\mathbf{x}_0\|_1}{\sqrt{k}\|\mathbf{x}_0\|_2} \|\mathbf{E}_1\|_F. \quad (\text{B.11})$$

■

C Results on non-convex recovery

Next two lemmas are standard results on Gaussian measurement operators.

Lemma C.1 (Properties of Gaussian mappings) Assume \mathbf{X} is an arbitrary matrix with unit Frobenius norm. An i.i.d. Gaussian measurement operator $\mathcal{G}(\cdot)$ satisfies the following:

- $\mathbb{E}[\|\mathcal{G}(\mathbf{X})\|_2^2] = m$.
- There exists an absolute constant $c > 0$ such that, for all $1 \geq \varepsilon \geq 0$, we have

$$\mathbb{P}(|\|\mathcal{G}(\mathbf{X})\|_2^2 - m| \geq \varepsilon m) \leq 2 \exp(-c\varepsilon^2 m). \quad (\text{C.1})$$

Proof. Observe that, when $\|\mathbf{X}\|_F = 1$, entries of $\mathcal{G}(\mathbf{X})$ are i.i.d. standard normal. Hence, the first statement follows directly. For the second statement, we use the fact that square of a Gaussian random variable is sub-exponential and view $\|\mathcal{G}(\mathbf{X})\|_2^2$ as a sum of m i.i.d. subexponentials. Then, result follows from Corollary 5.17 of [50].

■

For the consequent lemmas, \mathcal{S} denotes the unit Frobenius norm sphere in $\mathbb{R}^{n_1 \times n_2}$.

Lemma C.2 Let $\mathcal{D} \in \mathbb{R}^{n_1 \times n_2}$ be an arbitrary cone and $\mathcal{G}(\cdot) : \mathbb{R}^{n_1 \times n_2} \rightarrow \mathbb{R}^m$ be an i.i.d. Gaussian measurement operator. Assume that the set $\bar{\mathcal{D}} = \mathcal{S} \cap \mathcal{D}$ has ε -covering number bounded above by $\eta(\varepsilon)$. Then, there exists constants $c_1, c_2 > 0$ such that whenever $m \geq c_1 \log \eta(1/4)$, with probability $1 - 2\exp(-c_2 m)$, we have

$$\mathcal{D} \cap \mathcal{N}(\mathcal{G}) = \{0\}.$$

Proof. Let $\eta = \eta(\frac{1}{4})$, and $\{\mathbf{X}_i\}_{i=1}^\eta$ be a $\frac{1}{4}$ -covering of $\bar{\mathcal{D}}$. With probability at least $1 - 2\eta \exp(-c\varepsilon^2 m)$, for all i , we have

$$(1 - \varepsilon)m \leq \|\mathcal{A}(\mathbf{X}_i)\|_2^2 \leq (1 + \varepsilon)m. \quad (\text{C.2})$$

Now, let $\mathbf{X}_{\text{sup}} = \arg \sup_{\mathbf{X} \in \bar{\mathcal{D}}} \|\mathcal{A}(\mathbf{X})\|_2$. Choose $1 \leq a \leq \eta$ such that $\|\mathbf{X}_a - \mathbf{X}_{\text{sup}}\|_2 \leq 1/4$. Then:

$$\|\mathcal{A}(\mathbf{X}_{\text{sup}})\|_2 \leq \|\mathcal{A}(\mathbf{X}_a)\|_2 + \|\mathcal{A}(\mathbf{X}_{\text{sup}} - \mathbf{X}_a)\|_2 \leq (1 + \varepsilon)m + \frac{1}{4}\|\mathcal{A}(\mathbf{X}_{\text{sup}})\|_2. \quad (\text{C.3})$$

Hence, $\|\mathcal{A}(\mathbf{X}_{\text{sup}})\|_2 \leq \frac{4}{3}(1 + \varepsilon)m$. Similarly, let $\mathbf{X}_{\text{inf}} = \arg \inf_{\mathbf{X} \in \bar{\mathcal{D}}} \|\mathcal{A}(\mathbf{X})\|_2$. Choose $1 \leq b \leq \eta$ satisfying $\|\mathbf{X}_b - \mathbf{X}_{\text{inf}}\|_2 \leq 1/4$. Then,

$$\|\mathcal{A}(\mathbf{X}_{\text{inf}})\|_2 \geq \|\mathcal{A}(\mathbf{X}_b)\|_2 - \|\mathcal{A}(\mathbf{X}_{\text{inf}} - \mathbf{X}_b)\|_2 \geq (1 - \varepsilon)m - \frac{1}{3}(1 + \varepsilon)m. \quad (\text{C.4})$$

This yields $\|\mathcal{A}(\mathbf{X}_{\text{inf}})\|_2 \geq \frac{2-4\varepsilon}{3}m$. Choosing $\varepsilon = 1/4$ whenever $m \geq \frac{32}{c} \log(\eta)$ with the desired probability, $\|\mathcal{A}(\mathbf{X}_{\text{inf}})\|_2 > 0$. Equivalently, $\bar{\mathcal{D}} \cap \mathcal{N}(\mathcal{A}) = \emptyset$. Since $\mathcal{A}(\cdot)$ is linear and \mathcal{D} is a cone, the claim is proved. ■

The following lemma gives a covering number of the set of low rank matrices.

Lemma C.3 (Candes and Plan, [10]) Let M be the set of matrices in $\mathbb{R}^{n_1 \times n_2}$ with rank at most r . Then, for any $\varepsilon > 0$, there exists a covering of $\mathcal{S} \cap M$ with size at most $(\frac{c_3}{\varepsilon})^{(n_1+n_2)r}$ where c_3 is an absolute constant. In particular, $\log(\eta(1/4))$ is upper bounded by $C^{(n_1+n_2)r}$ for some constant $C > 0$.

Now, we use Lemma C.3 to find the covering number of the set of simultaneously low rank and sparse matrices.

C.1 Proof of Lemma 5.2

Proof. Assume S has $\frac{1}{4}$ -covering number N . Then, using Lemma C.2, whenever $m \geq c_1 \log N$, (5.1) will hold. What remains is to find N . To do this, we cover each individual $d_1 \times d_2$ submatrix and then take the union of the covers. For a fixed submatrix, using Lemma C.3, $\frac{1}{4}$ -covering number is given by $C^{(d_1+d_2)s}$. In total there are $\binom{n_1}{d_1} \times \binom{n_2}{d_2}$ distinct submatrices. Consequently, by using $\log \binom{n}{d} \approx d \log \frac{n}{d} + d$, we find

$$\log N \leq \log \left(\binom{n_1}{d_1} \times \binom{n_2}{d_2} C^{(d_1+d_2)s} \right) \leq d_1 \log \frac{n_1}{d_1} + d_1 + d_2 \log \frac{n_2}{d_2} + d_2 + (d_1 + d_2)s \log C,$$

and obtain the desired result. ■